PdfPlus User Manual

Innovative Technology

PdfPlus User Manual

Version 1b424e3 (master), 11 April 2013
Copyright © 2013 William Matthew Limited. All rights reserved.

Acknowledgements.

Microsoft and Windows are registered to Microsoft Corporation.
HP and PCL are trademarks of Hewlett Packard.
PowerEdge is a trademark of Dell Corporation.
VMware is a trademark of VMware, Inc.
Linux is a registered trademark of Linus Torvalds.
Prescribe is a registered trademark of Kyocera-Mita Corporation.
E&OE

Table of Contents

About this manual vii
1. Intended Audience vii
2. Software Licence and Copyright vii
1. Introduction 9
2. Getting Started 11
2.1. Platform set-up 11
Setting up a Tempest PdfPlus 11
Setting up a Server-based PdfPlus 12
Setting up a Virtualized PdfPlus 13
2.2. Accessing the console 49
Accessing the console on Tempest appliances 49
Accessing the console on Virtualized appliances 49
Accessing the console on Server-based appliances 50
2.3. IP address configuration 50
Console IP address configuration 50
2.4. Accessing the PdfPlus user interface 51
PdfPlus users 51
3. Disaster recovery and backup procedures 53
3.1. Backup 53
Automated backup with WML's disaster recovery system 53
Static virtual machine backup procedure 53
Database backups 55
Making backups of PdfPlus system settings 55
3.2. Recovery 55
4. PdfPlus Configuration 57
4.1. System configuration 57
Network settings 57
System date and time 59
System timezone and locale 60
System logging 60
Email settings 63
Licence key 64
System passwords 65
System settings 66
Firmware 68
Status and Control 69
5. Queue Management 71
5.1. Queue management interface 71
Print user actions 72
Print admin controls 73
Classes 75
5.2. Queue configuration 76
Data input 76

Overlay templates and logic profiles 83
Archiving 84
Failover 87
Alerts (Queue Monitoring) 87
PPD driver 89
Driver options 91
Output 92
Queue info 94
5.3. Queue configuration: Example workflow 94
5.4. Jobs 99
5.5. Platform settings 99
Cups Control 99
Failover 99
Edit cupsd.conf 100
Edit mime.types 100
Edit mime.convs 100
6. Template Management 101
6.1. Introduction 101
6.2. Overlay elements 103
6.3. Overlay templates 104
6.4. Overlay template list 104
6.5. Overlay template manager 106
Saving changes to the template 107
Template management 107
Template settings 108
The canvas 112
The overlay elements user interface 115
Testing print data 125
6.6. Overlay Templates: Example workflow 126
Create a print queue 126
Create PDF forms 126
Create and name a new template 127
Upload and activate PDF forms 128
Add text overlays 130
Review template settings 132
Upload and test 132
Bind the template to the queue 133
6.7. Using structured input data 134
XML input data 135
CSV data 140
Jetforms format input data 141
Template special keys 142
6.8. Using data groups 142
Letter with variable content 142
Multiple page invoice 147

7. Logic Profiles 157
8. File Manager 163
8.1. Form Files 164
8.2. Log Files 164
8.3. Archive Files 164
8.4. PPD Files 164
8.5. All User Files 164
Print Diagnostics 164
9. Print Reports 165
9.1. Viewing reports in the web interface 167
9.2. Reports with a pair of DataGateway devices 168
10. Kyocera Prescribe Emulation 171
10.1. Basic Defaults 172
10.2. Paper Cassettes 172
Using Digital Paper Stock 173
Using PostScript Paper Selection Tags 173
10.3. Persistent State 173
11. Appliance Manager 175
11.1. The Appliance Manager push 175
11.2. The user interface 175
11.3. Appliances and groups 176
Populating the appliance list 177
Appliance status 178
Appliance report 179
Managing pushes 181
11.4. Stored configurations 183
11.5. Stored firmware 184
11.6. Appliance Manager for Disaster Recovery 184
11.7. Appliance Manager Networking 185
12. WML Secure Forms 187

About this manual

This manual describes the use of WML PdfPlus software.

1. Intended Audience

This manual is aimed at IT administration staff who will be involved with maintaining and configuring the PdfPlus and at print production staff who will be reviewing reports.

2. Software Licence and Copyright

The Software is copyrighted and all rights reserved by William Matthew Limited, 22 Haxey Lane, Haxey, Doncaster, South Yorkshire, DN9 2NE. UK. Email Info@wmltd.co.uk

> **Important**
>
> Before using the WML PdfPlus Software please read the licence agreement carefully. The right to use this software is granted only if the Customer agrees to the terms of this Licence. If you do not agree to the terms of this Licence, please return the product unused.
>
> William Matthew Limited considers that installation of the Software is your acceptance of our terms and conditions of the license.

1. The Software has been developed for use with the WML PdfPlus; The Licence is to be used without any time limit and may be subject to change.
2. The Software may not be duplicated, copied, resold, or re-branded without the written consent of William Matthew Limited.
3. It is the responsibility of any user, reseller, distributor, OEM or other third party to protect the rights of William Matthew Limited against violation of this Software Licence.
4. Some portions of The Software have been developed by third parties and are used by WML under the terms of the GNU General Public Licence, as well as other open source and free software licences. The source code to these portions of The Software are available to WML's customers under the same terms under which WML uses them. WML will uphold and protect the copyright of the original author of these portions of The Software.
5. William Matthew will consider any attempt to duplicate the functionality of any application provided by the software to constitute a breach of software licence, and may instigate legal proceedings to protect the copyright of the PdfPlus software and William Matthew Limited.
6. This Licence and the Software product may be transferred to a third party provided the third party agrees to all the terms of this Licence and that the transferring party retains no copies of the Software.
7. William Matthew Limited reserves the right to terminate this Licence upon breach.

8. In the event of a termination, all copies of the software must be destroyed or returned to William Matthew Ltd
9. Any disputes in respect of the Software Licence will be considered a breach of copyright and subject to English Law.

Chapter 1. Introduction

WML PdfPlus is a network appliance for workgroup printing. It acts as a print server on your network, receiving data from a variety of sources, which it can manipulate and pass on to your printers.

PdfPlus can accept data from your mainframe, Unix, Windows and Linux application servers using Berkeley LPD protocol, IPP protocol and simple TCP raw-data transfer. It is also able to collect print jobs from Windows (SMB/CIFS) and Unix (NFS) network-shared "hot folders". In addition, FTP can be used to send input data, enabling the PdfPlus to be used in conjunction with multi-function printers to process scanned images.

PdfPlus can work with pre-formatted print data or with structured-text data such as XML and comma separated value data. Acceptable pre-formatted print data types are PDF, Postscript and formatted plain text. Structured-text input data can be formatted in XML or as comma separated values (CSV).

PdfPlus is configured via its built in web interface. It provides a graphical *overlay template editor* which allows you to build up a template comprised of PDF forms and *overlay elements* such as the date, time, barcodes and static text.

A *logic engine* allows the PdfPlus to trigger overlay templates based on the content of incoming print data.

Available as an optional module is the *Secure Forms* application for retail environments. Secure Forms reveals a customer interface to the overlay template which allows the customer to enter their details into a graphical form using a tablet computer.

PdfPlus can masquerade as multiple printers by assuming an IP address for each printer. Using this feature, it is possible to use PdfPlus to replace a number of desktop printers with a single production printer requiring no change to the application which generates the print jobs.

Print jobs are output, via a print driver, to destination printers. LPD, IPP or raw protocols can be used. Archiving of jobs (in PDF format) is possible.

PdfPlus has the versatility to help you to implement new print functionality or to make a change in printer hardware with minimal change to your existing software.

Chapter 2. Getting Started

In this chapter, you will be guided through the initial set-up for the WML PdfPlus. PdfPlus is available in three platform configurations:

- *Tempest PdfPlus* is an appliance form-factor device, based on an AMD Geode LX processor.
- *Server-based PdfPlus* is a Dell or Intel server appliance, for situations where faster processing or local (hard disk) storage is required. Rack servers provide more network interfaces than the Tempest platform, which can be a requirement in enterprise networks.
- *Virtualized PdfPlus* is a VMware virtual machine image of the software. Virtualized PdfPlus can be delivered as a pure image to install on the customer's own VMware infrastructure, or on server hardware delivered and maintained by WML. Virtual print appliances are also available in a format suitable for running under Linux KVM (as used in Ubuntu Linux).

2.1. Platform set-up

In this section, you will find out how to set up the various hardware and virtualized platform configurations. The following sections describe set-up tasks that are common to all configurations.

Setting up a Tempest PdfPlus

Figure 2.1. Tempest PdfPlus appliance and stand

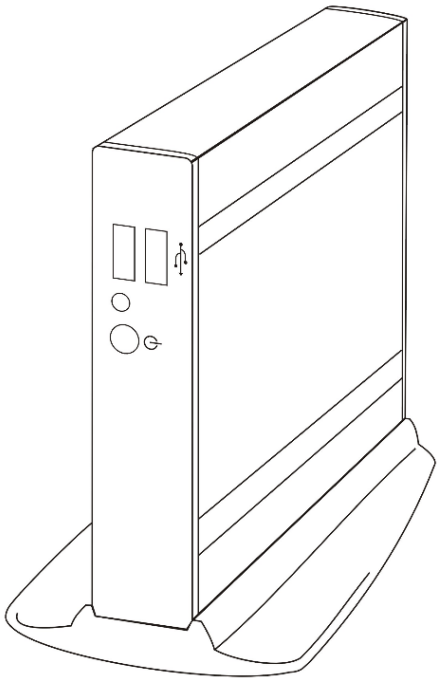

The Tempest PdfPlus includes the following components:

- Tempest appliance;
- stand;
- power adaptor.

You will also need an ethernet cable to connect the unit to your network.

Figure 2.2. Rear panel of Tempest PdfPlus

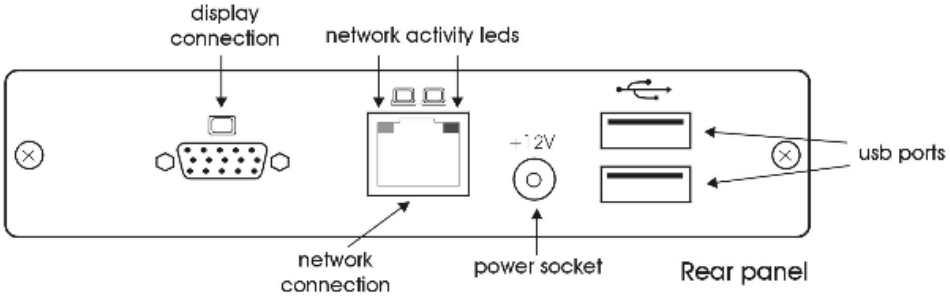

Connect the box to your network by plugging an ethernet cable into the network port on the back of the unit (Figure 2.2).

Figure 2.3. Front panel of Tempest PdfPlus

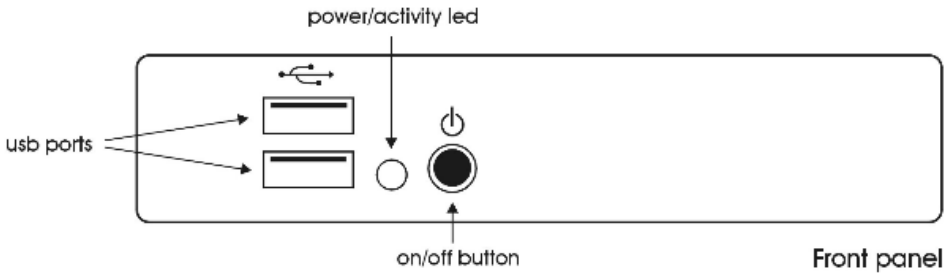

Plug the power supply into the back of the unit. Turn the PdfPlus on using the button on the front of the unit (Figure 2.3). You should see the blue indicator light come on.

Setting up a Server-based PdfPlus

> **Important**
>
> The exact choice of server equipment will vary depending on the expected workload for PdfPlus. For this reason, it's not possible to give a comprehensive setup guide for a Server-based PdfPlus here.

The Server-based PdfPlus will consist of a one or two unit rack-mount server, with a suitable number of processors, RAM and disk storage for the intended application and print throughput. It will be delivered as a single unit, complete with a rack-mount kit.

The Server-based PdfPlus will be delivered fully configured and ready to run. Power the device on and direct your web browser to the IP address provided by WML staff.

Setting up a Virtualized PdfPlus

Virtualized appliance disk files are distributed either as VMWare images (.vmdk files) or a Linux Kernel Virtual Machine (KVM) compatible images (.qcow2 files).

Follow the procedure given by the vendor of your virtualization infrastructure to create a new virtual machine using the Virtualized appliance disk file as the virtual disk.

> **Note**
>
> The Virtualized appliance disk file will be read from and written to. For this reason, it is advisable to keep a copy of the original disk file from which a fresh virtual machine can be created. Select a suitable location on your disk system to store your in-use disk files. This should be somewhere where it is unlikely that the files will be moved or interfered with externally.
>
> Make a separate copy of the Virtualized appliance disk file for each new virtual appliance which you create.

Adding a second virtual disk for increased storage

The Virtualized appliance disk file provides up to 2.5 GB of space to store settings, DCC profiles and PDF forms. This is a sensible amount of storage for systems which do not need to archive locally, store log files locally for long periods of operation or run an internal database.

In practice, DataGateway systems typically require more storage as they will, in most cases, run an internal database (MySQL), archive document data and store log files internally. PdfPlus and PrePrint systems usually work well with the default disk file.

All WML Print Platform appliances will recognise the presence of a second virtual hard disk and make use of this disk to store all persistent data in place of the persistent partition of the Virtualized appliance disk file.

When the second disk is first detected during a boot of the Virtualized appliance, the system tests it to see if it already contains the persistent files from a WML Print Platform system. If it does contain files, it is attached as the persistent storage, and boot up continues as normal. If it is empty, then the boot up procedure first copies all the files from the Virtualized appliance disk file's persistent partition to the second disk, then completes boot up using the second disk as persistent storage.

To make use of a second disk, simply create a new disk in VMware or KVM, create a Linux ext3 filesystem on the entire disk (that is, don't create any partitons on the new disk), ensure it is attached to your virtual machine as the second disk and reboot your appliance.

Ideally, the new disk should appear as a SCSI disk to the operating system. It should appear internally on the Linux-based WML Print Platform system as the device **/dev/sdb** for correct results.

VMware infrastructure options

VMware provides a number of different infrastructure options, including VMware workstation, VMware ESX and VMware vSphere/ESXi.

A Virtualized PdfPlus should work on any type of VMware infrastructure, from the free VMware Server, through to VMware vSphere systems, however, you do need to choose to download the correct vmdk disk image from our servers.

VMware virtual disks, pre-installed with PdfPlus software are available in two formats. The first is a growable disk provided as a single file with a name such as:

```
wmlpp-pplus-vmware-sata-tree3-20130228_1502.vmdk
```

This format can be used as the virtual disk for virtual machines created on VMware Workstation and VMware Server, but not recent versions of VMware ESXi/vSphere.

If you are using VMware ESXi or vSphere, you will need to download the ESX-compatible VMDK. This is provided as a zip file with a file name suffix **.esx.zip**. The zip contains two files, a .esx.vmdk disk descriptor file and a 2.5 GB pre-allocated .esx-flat.vmdk file which contains the disk data. For example:

```
wmlpp-pplus-vmware-sata-tree3-20130228_1502.esx.vmdk
wmlpp-pplus-vmware-sata-tree3-20130228_1502.esx-flat.vmdk
```

Upload both of these .vmdk files to the data store of your VMware infrastructure before creating a new virtual machine.

> **Note**
>
> It is possible to use the ESX-compatible VMDK files on VMware Workstation.

VMware Server walk-through

This walk-through demonstrates the set up of a Virtualized appliance on VMware Server version 2.0.2.

> **Note**
>
> It is assumed that you have VMware Server version 2.0.2 already installed on your PC. For help installing VMware Server, please refer to the VMware website.
>
> VMware Server 2.0.2 can be obtained both for Linux and for Windows platforms; our appliances have been verified to work correctly on both.

Creating a new Virtualized appliance is carried out in the VMware Infrastructure Web Access interface. This is accessed via a web browser on the system on which VMware Server was installed. Typically, it can be accessed at the address **http://yourcomputer:8333/**

Figure 2.4. Creating a Virtual Machine in the VMware Infrastructure Web Access Interface.

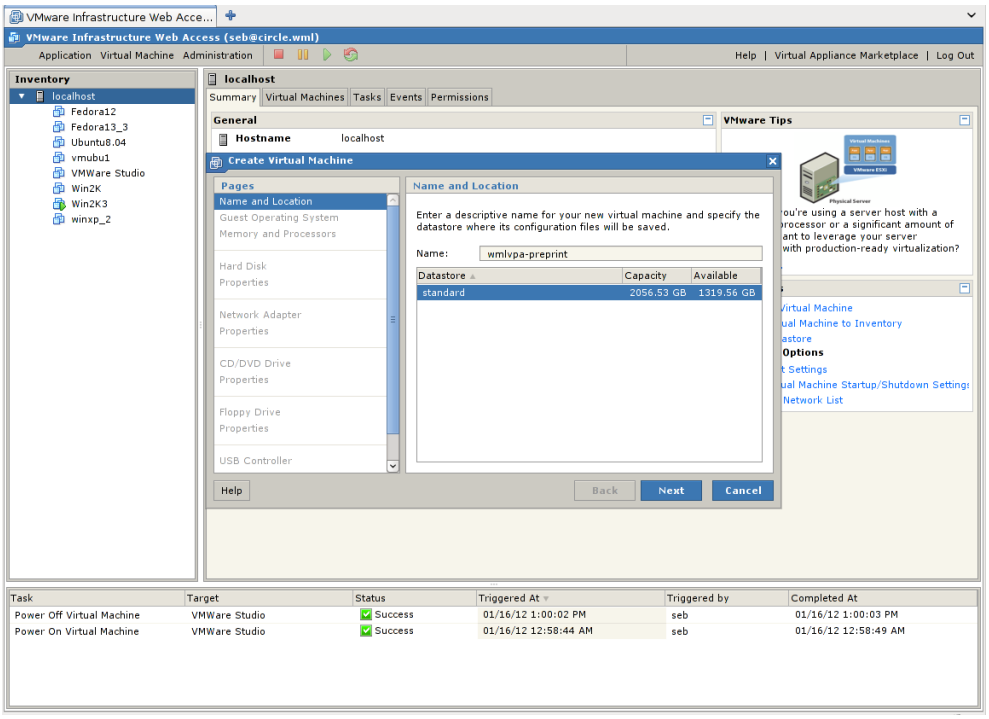

In VMware Infrastructure Web Access, click **Virtual Machine > Create Virtual Machine**. The **Create Virtual Machine** dialog will open (Figure 2.4). Assign a name for your new appliance. Here is has been called “wmlvpa-preprint”. Click **Next**.

Figure 2.5. Specifying the operating system for the guest appliance.

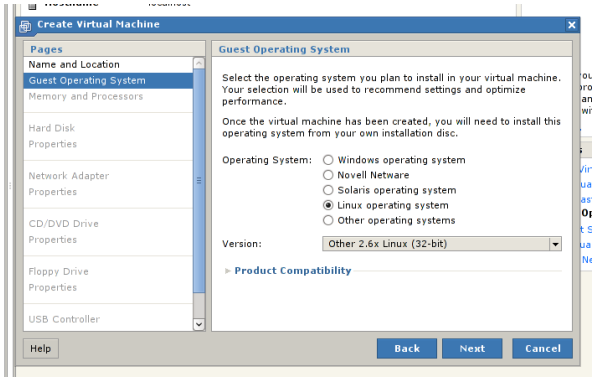

The next window (Figure 2.5) allows you to choose which operating system your target appliance will be running. All WML appliances run a 32 bit Linux operating system and so you should choose **Linux operating system** for **Operating System** and **Other 2.6x Linux (32-bit)** for **Version**. Click **Next**.

Figure 2.6. Setting the amount of RAM and number of processors for the virtual appliance.

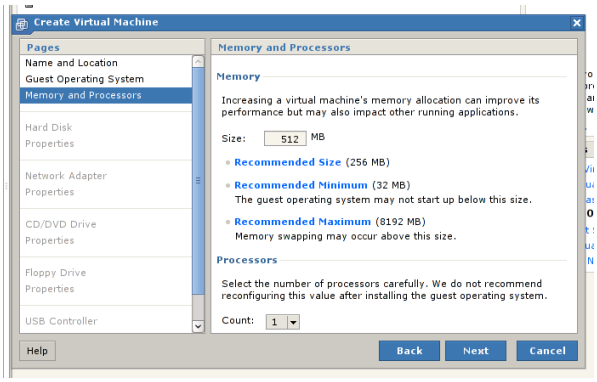

This screen (Figure 2.6) allows you to specify the initial amount of RAM and processor resource to assign to your virtual print appliance. Specify a minimum of 256 MB of RAM. Some applications may require more memory to run correctly, but 256 MB is sufficient for simpler tasks and short documents.

Assign the desired number of processors for the virtual appliance. A single processor may be sufficient for simpler applications. Assigning two processors to the system will improve the responsiveness of the web interface when the print processing load is high.

> **Note**
>
> The amount of RAM may be altered after the virtual appliance has been created. Altering the number of processors is not recommended by VMware, but it can be done. WML appliances will safely handle a change in the number of available processors.

Figure 2.7. Adding a hard disk to the virtual appliance.

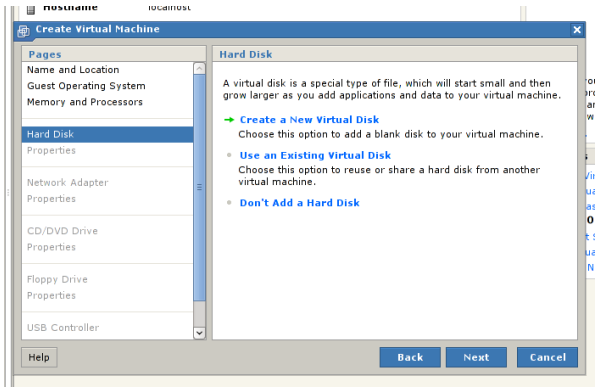

In the next window (Figure 2.7) you have the option to **Create a New Virtual Disk**. This would create an *empty* disk, which would then need to have the operating system installed (perhaps from a dvd or cdrom). Because you are using a pre-populated disk file supplied by WML, you should click on **Use an Existing Virtual Disk**, which will open the next window, shown in Figure 2.8.

Figure 2.8. Setting the hard disk properties.

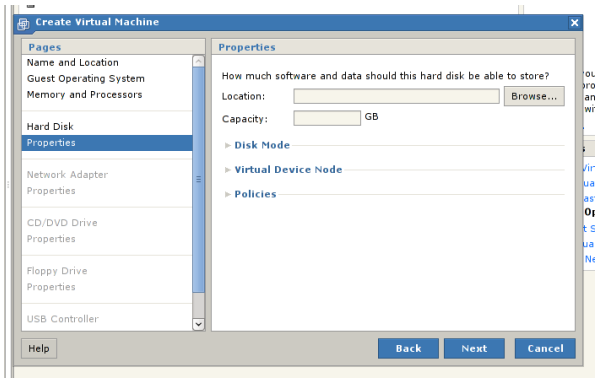

Figure 2.9. Selecting the WML Virtualized disk file.

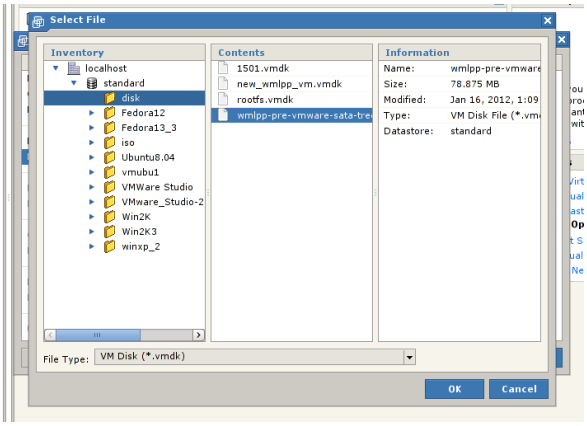

On the **Properties** window (Figure 2.8), click **Browse** and browse to the location where you saved your .vmdk disk image (See Figure 2.9) then click **OK** (it's not necessary to change **Disk Mode**, **Virtual Device Node** or **Policies** from their default values).

> **Warning**
>
> VMware Server will *not* make a copy of this disk image file; it's up to you to keep a copy of the original disk image. The disk image file will be read from and written to, and may increase in size to a maximum of approximately 2.5 GB.

Figure 2.10. The Network Adapter window.

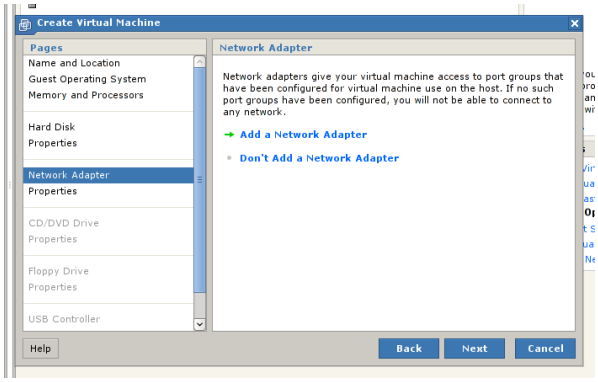

Figure 2.11. The Network Adapter properties window.

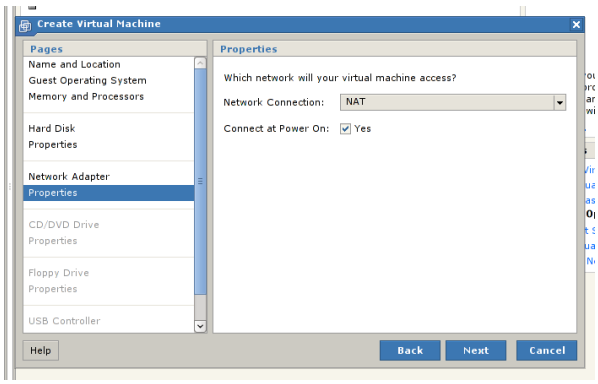

In the **Network Adapter** window (Figure 2.10), click **Add a Network Adapter**. In the properties window for the network adapter (Figure 2.11), you can select different kinds of network. For a most systems, you are likely to need **Bridged** networking, which allows your virtual appliance to join your physical network as if it were a real machine. **NAT** (Network Address Translation) networking is unlikely to be suitable, as the WML Virtualized appliance acts as a server. For test installations, **HostOnly** may be a suitable choice. In this case, the virtual machine exists on an internal, virtual network.

A full discussion of networking your virtual machines is outside the scope of this document. If you are unsure of which type of network connection to choose, contact your network administrator or your networking consultant or supplier.

Leave **Connect at Power On** checked, and press **Next**.

Figure 2.12. The CD/DVD Drive window.

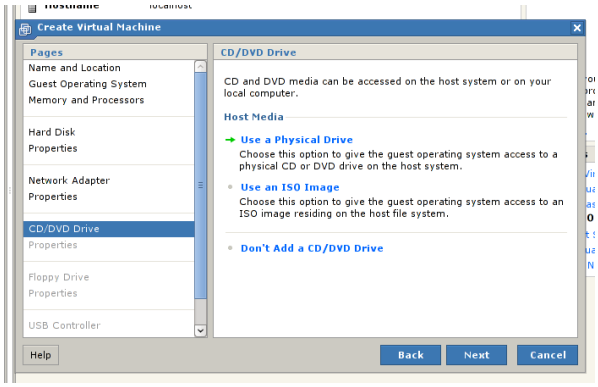

Figure 2.13. The Floppy Drive window.

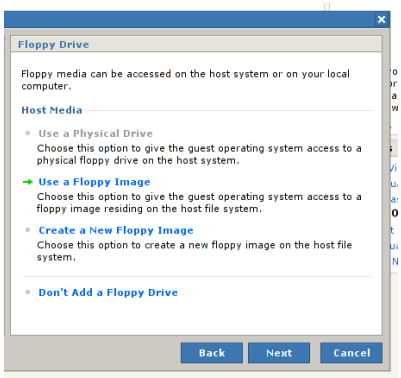

Figure 2.14. The USB Controller window.

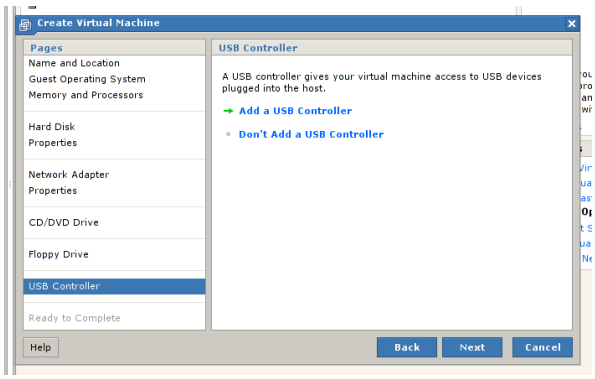

On the CD/DVD Drive window (Figure 2.12), click **Don't Add a CD/DVD Drive**. Likewise, on the floppy drive (Figure 2.13) and USB Controller (Figure 2.14) windows, click "Don't Add..."

Figure 2.15. Ready to Complete - a summary of your new virtual appliance.

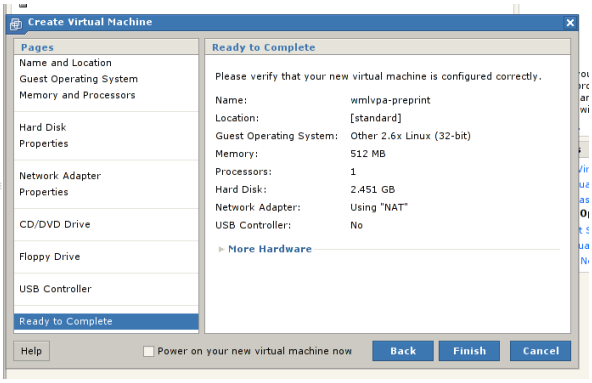

Click **Next** on the USB Controller window and you will arrive on the summary window (Figure 2.15). When you click **Finish**, your Virtualized virtual appliance is ready to use.

You can access the console of your virtual machines via the VMware Infrastructure Web Access interface. Figure 2.16 shows a Windows desktop with two running Virtualized appliances; a DataGateway and a PrePrint.

Figure 2.16. Two VMware virtual consoles; one is a WML DataGateway, the other is a WML PrePrint appliance.

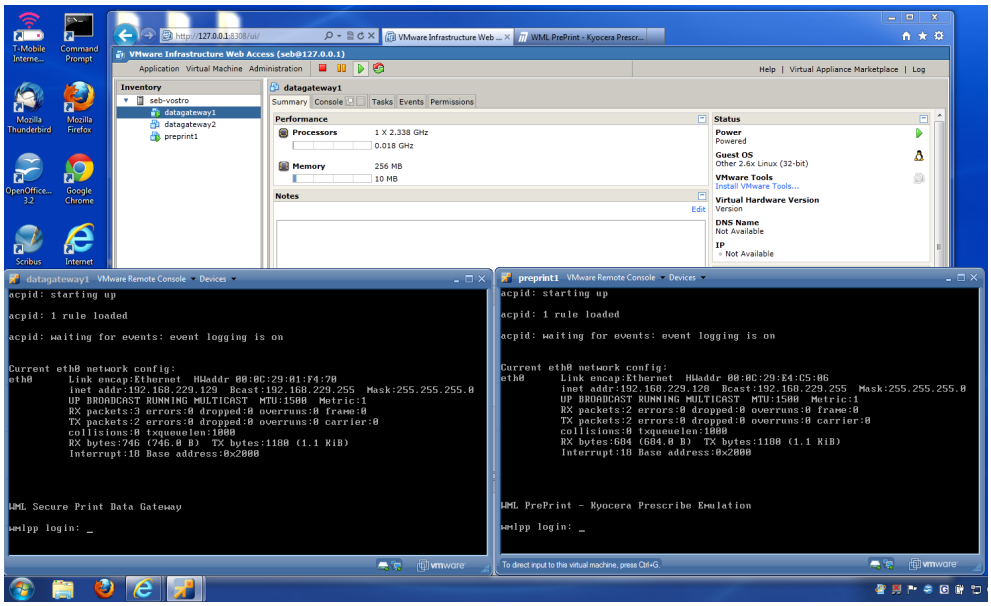

We hope this walk through for VMware Server is useful. Other VMware infrastructure options will differ from VMware Server, but the process will be similar to the one shown above.

VMware Workstation walk-through

This walk-through demonstrates the set up of a Virtualized appliance on VMware Workstation version 9.0 running on Microsoft Windows 8.

> **Note**
>
> It is assumed that you have VMware Workstation version 9.0 already installed on your PC. For help installing VMware Workstation, please refer to the VMware website.

Creating a new Virtualized appliance is carried out in the VMware Workstation user interface. This is accessed via an icon in your start menu or on the desktop.

Figure 2.17. The Home tab of the VMware Workstation user interface.

In the VMware Workstation Home tab, click **Create a New Virtual Machine**. The **New Virtual Machine Wizard** will open (Figure 2.18). Select **Custom (advanced)** so that it will be possible to specify that we want to use an existing, preformatted virtual disk. Click **Next**.

Figure 2.18. The "New Virtual Machine Wizard".

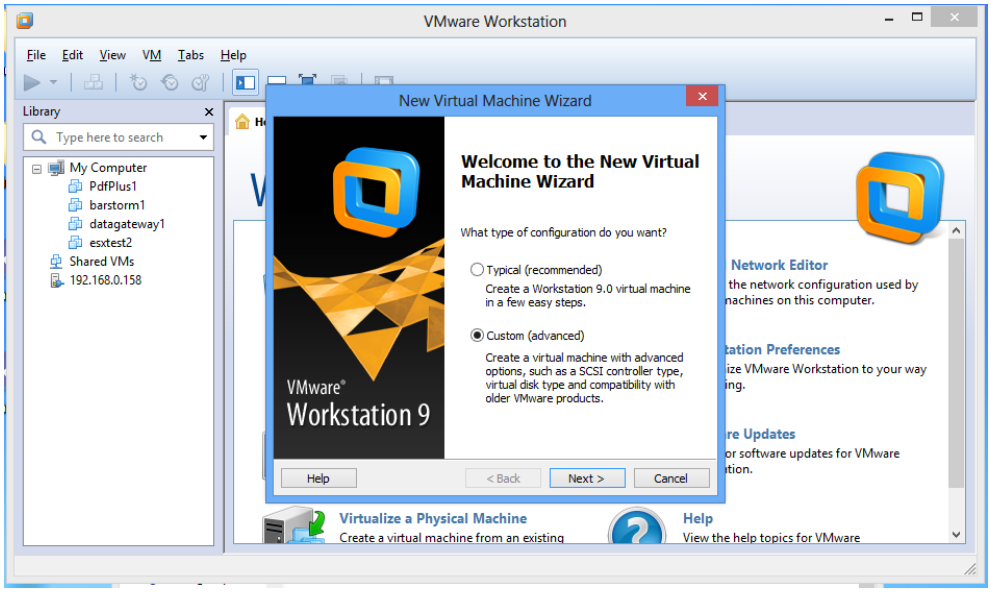

The first window allows you to specify the hardware compatibility for the virtual machine. This may be a consideration if you wish to move your virtual machine to older versions of VMware Workstation or to other VMware infrastructure products. In our example, we leave it at its default of **Workstation 9**.

Figure 2.19. Choosing the virtual machine hardware compatibility.

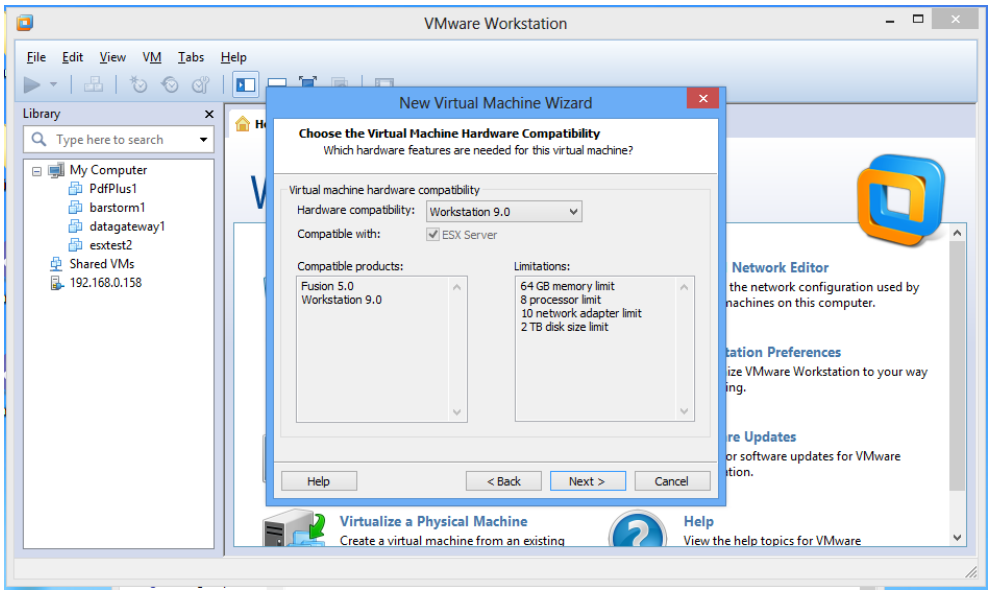

As we will use a pre-existing virtual disk, containing a ready-to-boot operating system, select **I will install the operating system later** on the **Guest operating system installation** window (Figure 2.20).

Figure 2.20. Selecting the source for the operating system for the virtual machine.

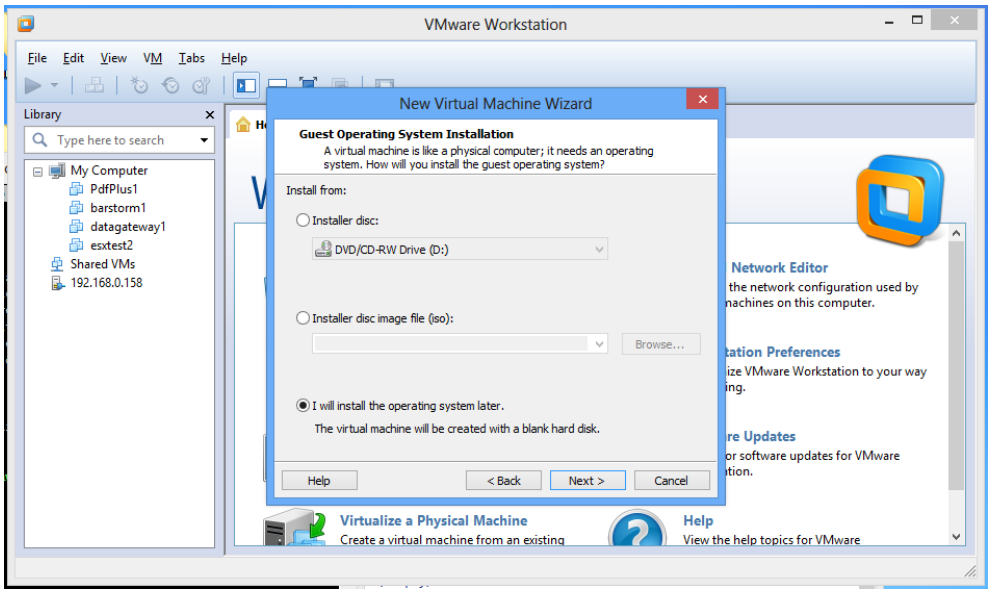

Figure 2.21. Specifying the operating system for the guest appliance.

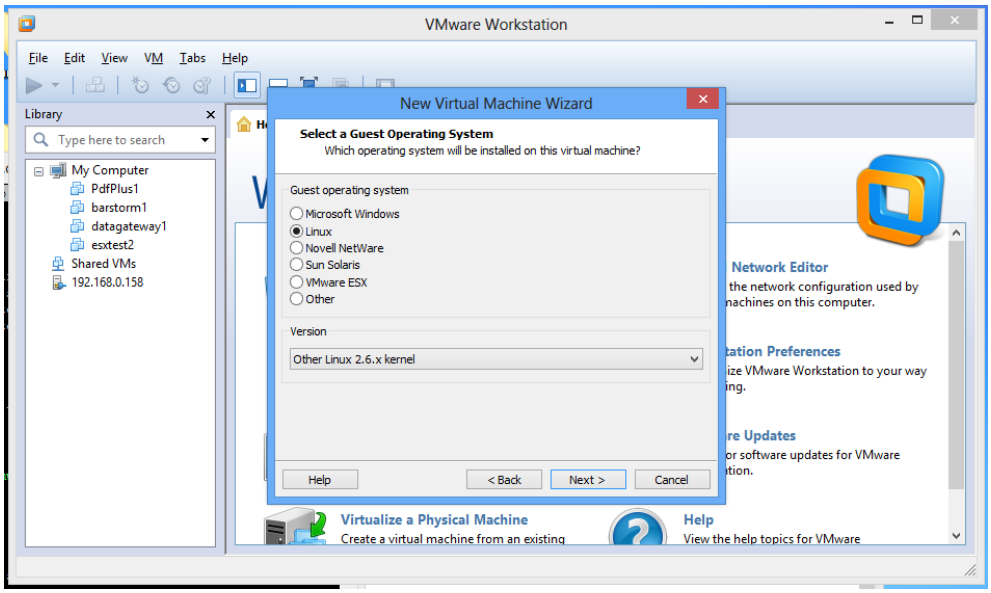

The next window (Figure 2.21) allows you to choose which operating system your target appliance will be running. All WML appliances run a 32 bit Linux operating system and so

you should choose **Linux** for **Guest operating system** and **Other Linux 2.6.x kernel** for **Version**. Click **Next**.

Figure 2.22. Choosing the name for the virtual machine.

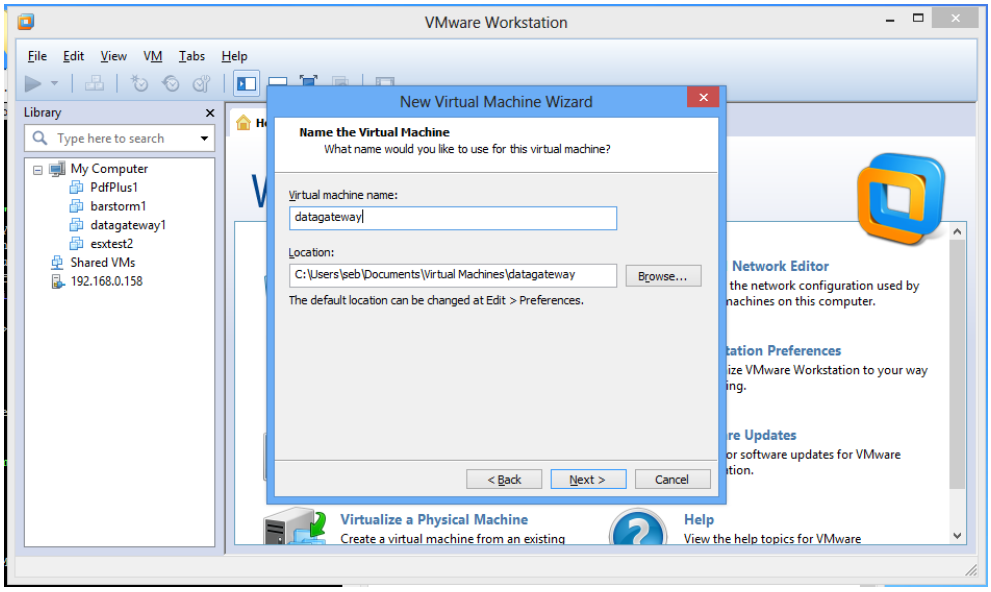

You can now choose the name for your virtual machine. Here, it has been named "datagateway" (Figure 2.22). The virtual machine definition files will be created in the default location **Documents\Virtual Machines**.

Figure 2.23. Setting the number of processors for the virtual appliance.

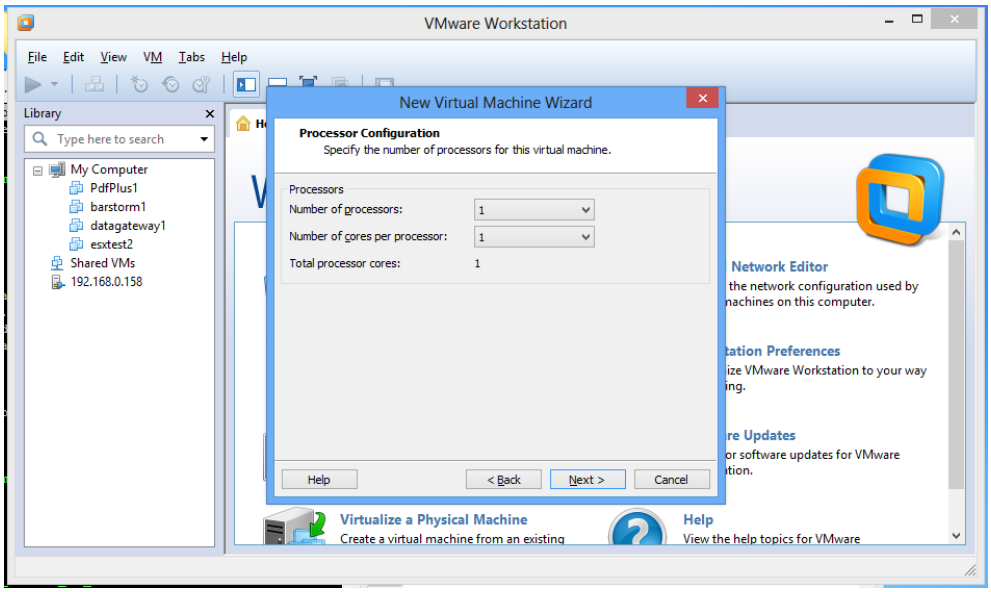

Figure 2.24. Setting the amount of RAM for the virtual appliance.

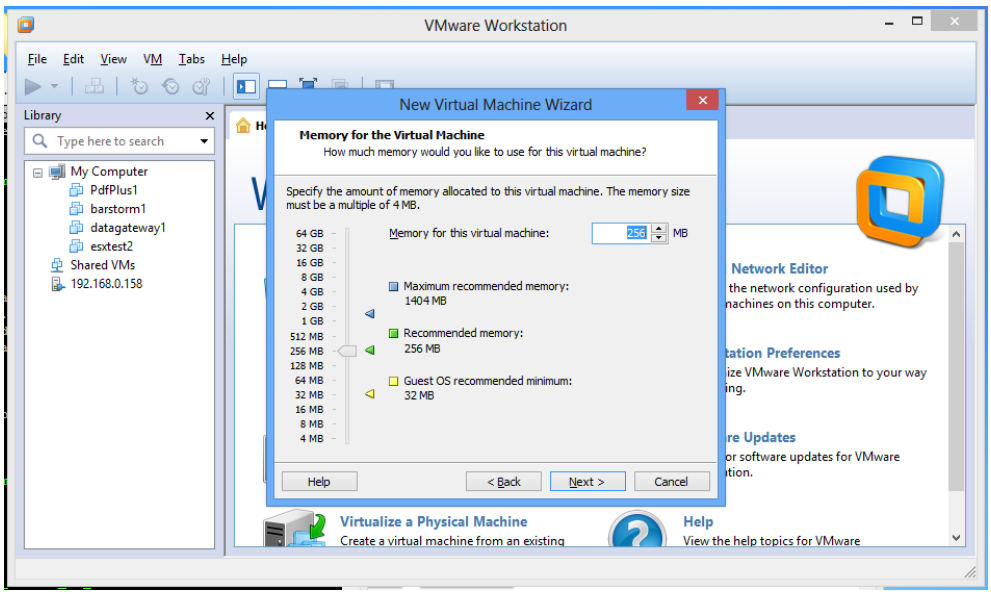

The next two screens (Figure 2.23 and Figure 2.24) allow you to specify the initial amount of RAM and processor resource to assign to your virtual print appliance.

Assign the desired number of processors for the virtual appliance. A single processor may be sufficient for simpler applications. Assigning two processors to the system will improve the responsiveness of the web interface when the print processing load is high.

Specify a minimum of 256 MB of RAM. Some applications may require more memory to run correctly, but 256 MB is sufficient for simpler tasks and short documents.

> **Note**
>
> The amount of RAM may be altered after the virtual appliance has been created. Altering the number of processors is not recommended by VMware, but it can be done. WML appliances will safely handle a change in the number of available processors.

Figure 2.25. The Network Type window.

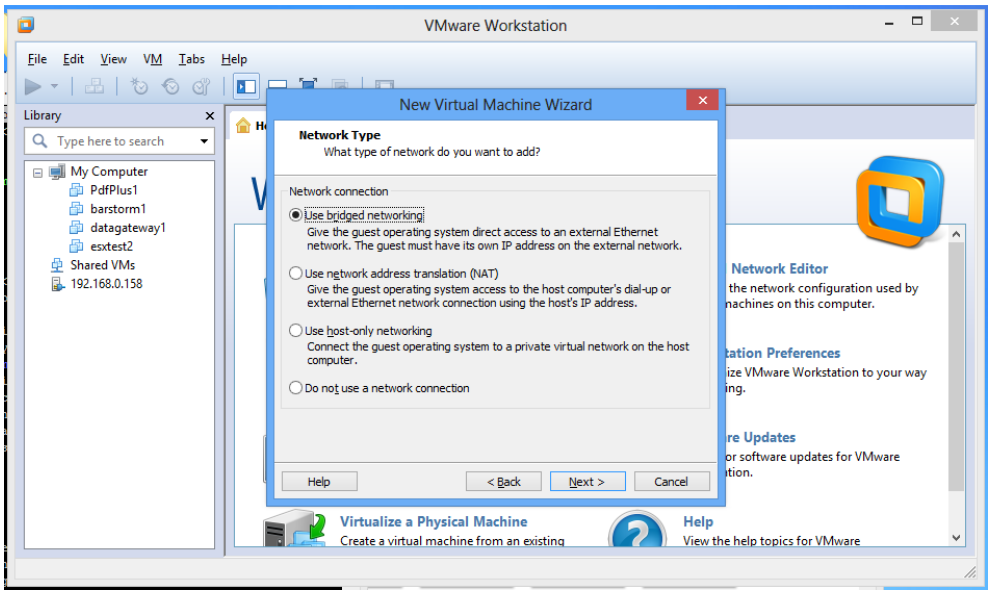

In the **Network Type** window (Figure 2.25) you can select different kinds of network. For a most systems, you are likely to need **Use bridged networking**, which allows your virtual appliance to join your physical network as if it were a real machine. **Use network address translation (NAT)** is unlikely to be suitable choice, as the WML Virtualized appliance acts as a server. For test installations, **Use host-only networking** may be a suitable choice. In this case, the virtual machine exists on an internal, virtual network.

A full discussion of networking your virtual machines is outside the scope of this document. If you are unsure of which type of network connection to choose, contact your network administrator or your networking consultant or supplier.

Figure 2.26. Selecting the virtual I/O controller type.

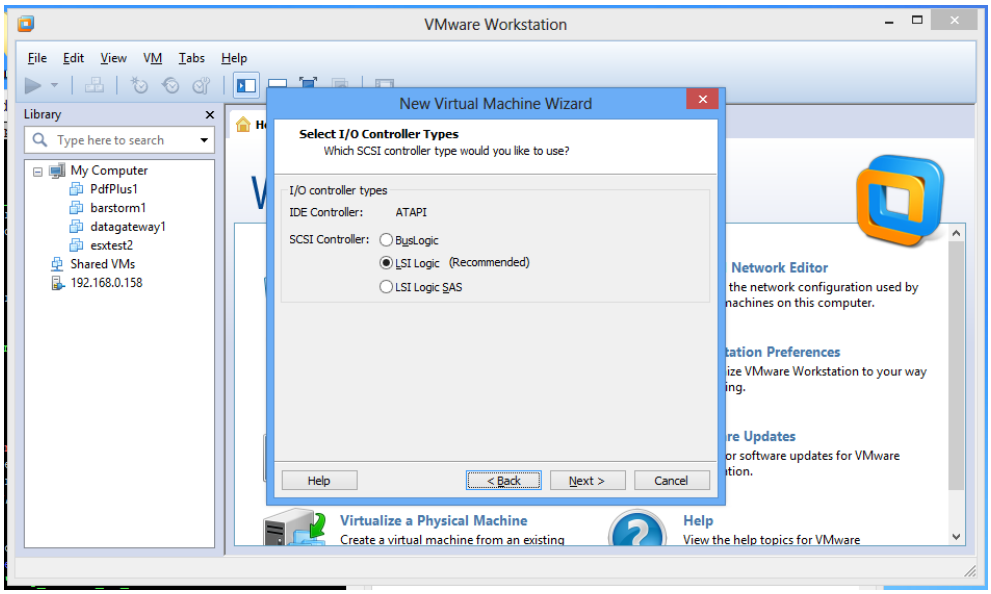

Figure 2.26 shows a screen allowing a choice of virtual disk I/O controllers. Leave the default **LSI Logic (Recommended)** checked.

Figure 2.27. Selecting a hard disk for the virtual appliance.

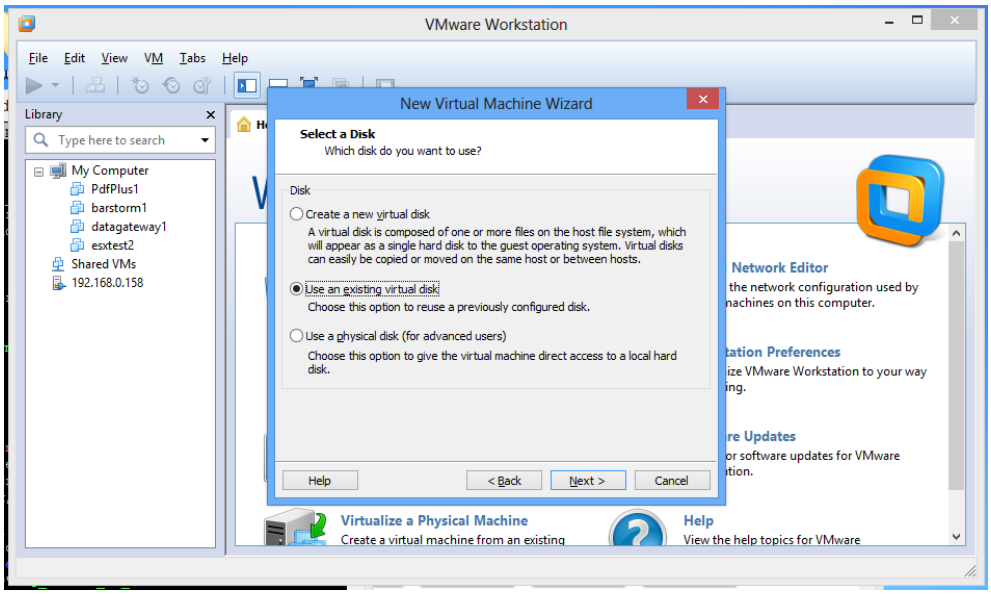

In the next window (Figure 2.27) you have the option to **Create a new virtual disk**. This would create an *empty* disk, which would then need to have the operating system installed (perhaps from a dvd or cdrom). Because you are using a pre-populated disk file supplied by WML, you should click on **Use an existing virtual disk**, which will open the next window, shown in Figure 2.28.

Figure 2.28. Selecting an existing disk.

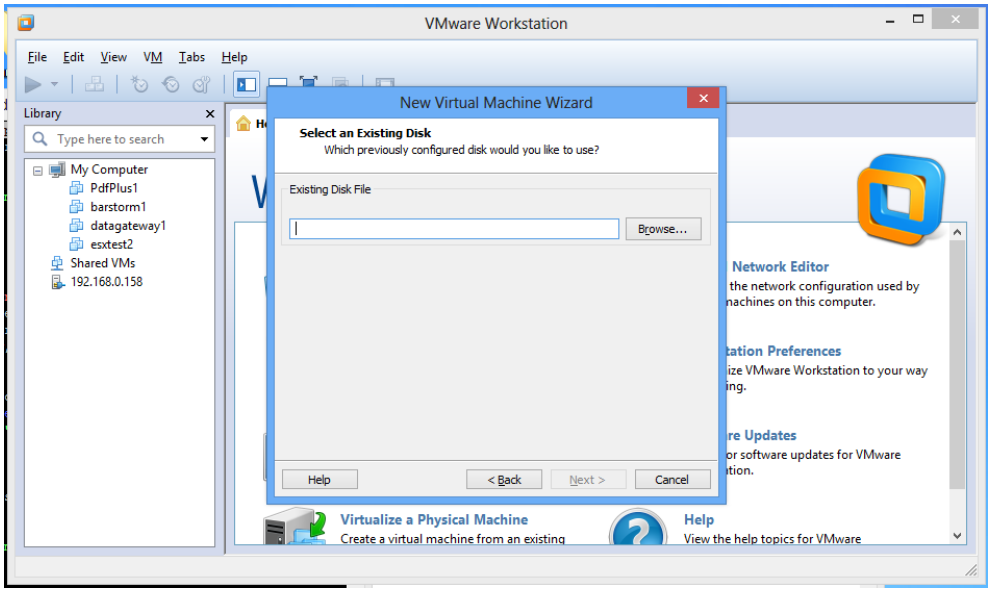

Browse to the vmdk disk image (Figure 2.29). You will probably see a pop-up asking if you wish to convert the existing virtual disk to a newer format (Figure 2.30). You can choose to do this or not, as you prefer; the virtual machine will boot correctly in either case.

> **Warning**
>
> VMware Workstation will *not* make a copy of this disk image file; it's up to you to keep a copy of the original disk image. The disk image file will be read from and written to, and may increase in size to a maximum of approximately 2.5 GB.

Figure 2.29. Selecting the WML Virtualized disk file.

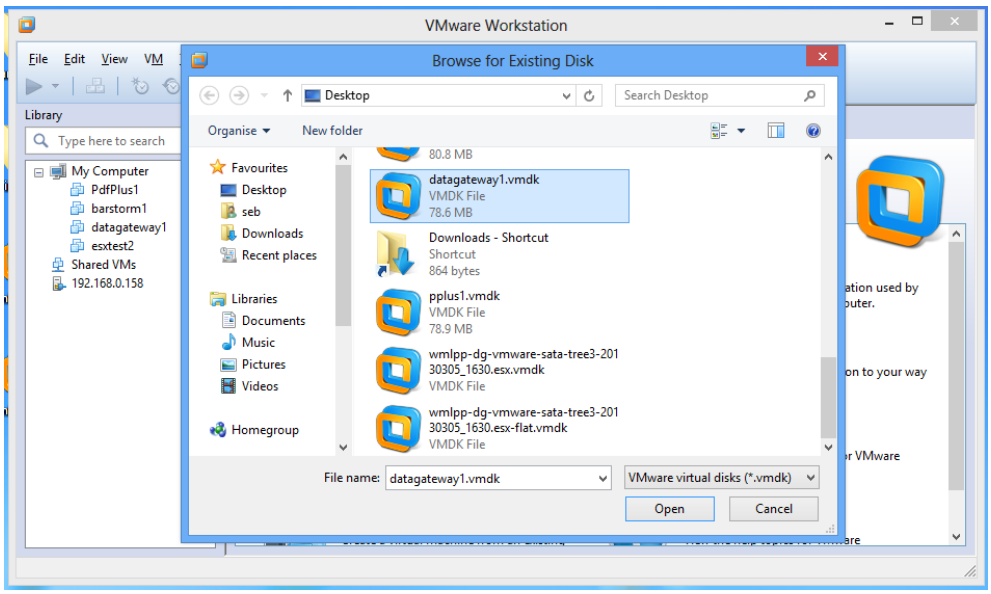

Figure 2.30. Converting to newer format.

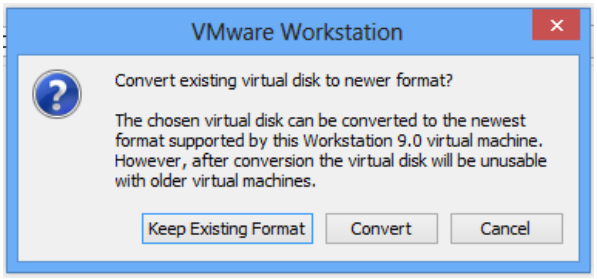

Choose to keep the existing disk format, or convert it and press **Next**. You'll see the final window in the wizard (Figure 2.31). When you click **Finish**, your Virtualized virtual appliance is ready to use.

Figure 2.31. Ready to Create Virtual Machine - a summary of your new virtual appliance.

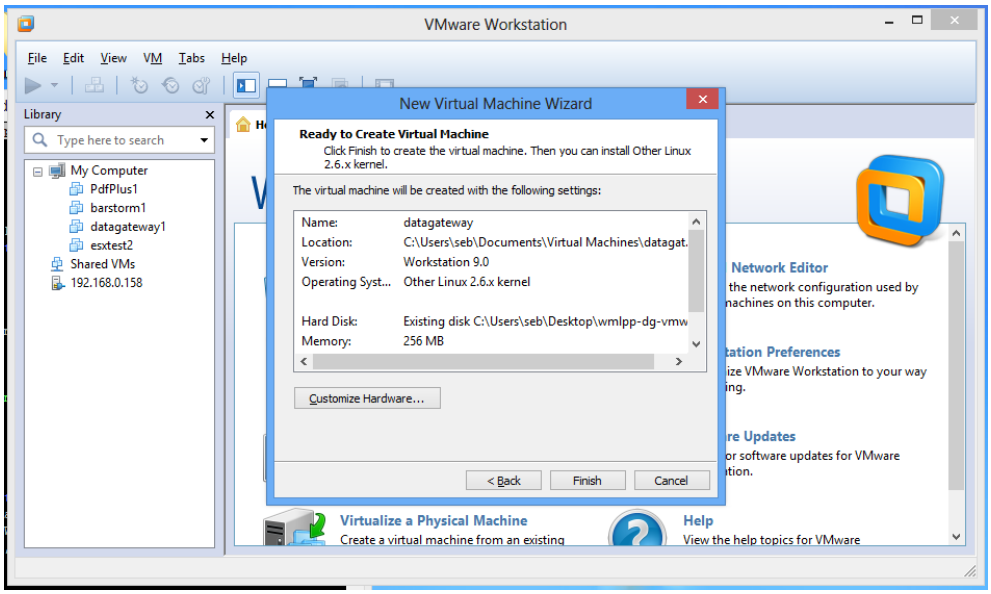

After creating your new virtual machine, a new tab will appear in the VMware Workstation UI. You can access the console of your virtual machine here (Figure 2.32). When you first boot your virtual machine, you will see a message about installing a guest operating system. You can simply click the button labelled **I Finished Installing**.

> **Note**
>
> There is no need to install VMware tools on a WML Virtualized. VMware tools enhances the user experience when using a mouse within a graphical environment. The PdfPlus has no graphical console and hence VMware tools is not used.

Figure 2.32. The new appliance tab.

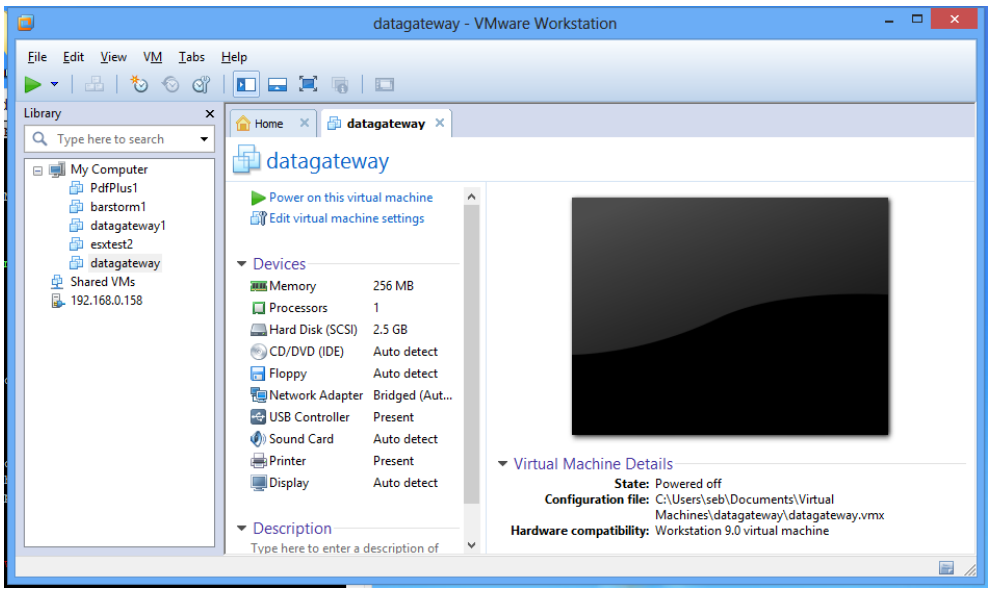

We hope you found this walk-through for VMware Workstation useful. If you compare this with the VMware Server walk-through, you will see the process has many common features.

VMware vSphere ESXi walk-through

This walk-through demonstrates the set up of a Virtualized appliance on a VMware ESXi hypervisor version 5.1.0. The ESXi hypervisor is controlled from vSphere client 5.1.0 running on Windows 8.

> **Note**
>
> It is assumed that you have a server with VMware ESXi version 5.1.0 already installed on a server, and VMware vSphere client version 5.1.0 on a separate PC or laptop. For help installing VMware vSphere, please refer to the VMware website.

Figure 2.33. The Summary tab in vSphere client for the ESXi host (192.168.0.158).

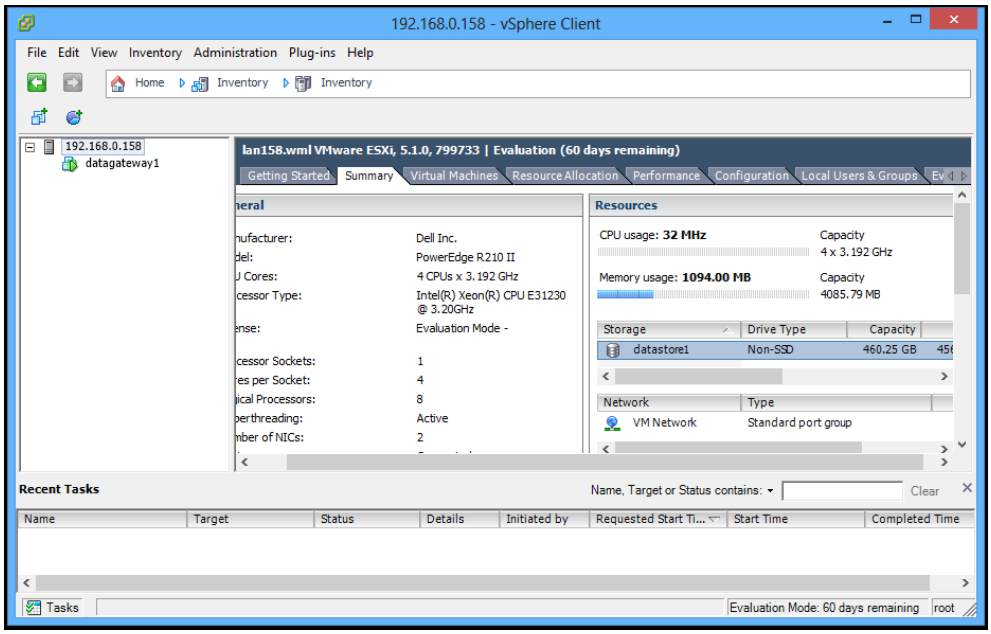

Creating a new Virtualized appliance is carried out in the VMware vSphere client user interface. This is accessed via an icon in your start menu or on the desktop. When you open vSphere client, you must connect to the IP address of your VMware ESXi host.

Once you have connected to the ESXi host, you will see the vSphere client **Getting Started** screen. Figure 2.33 shows the vSphere client, with the **Summary** tab showing.

Figure 2.34. VMware images on the WML web server.

Before starting, download the latest PdfPlus image from:

```
http://wmltd.co.uk/pdfplus/
```

Choose the .esx.zip file (Figure 2.34). Unpack the zip and find the two .vmdk files (Figure 2.35).

Figure 2.35. The .esx.zip file unpacks to provide two vmdk files.

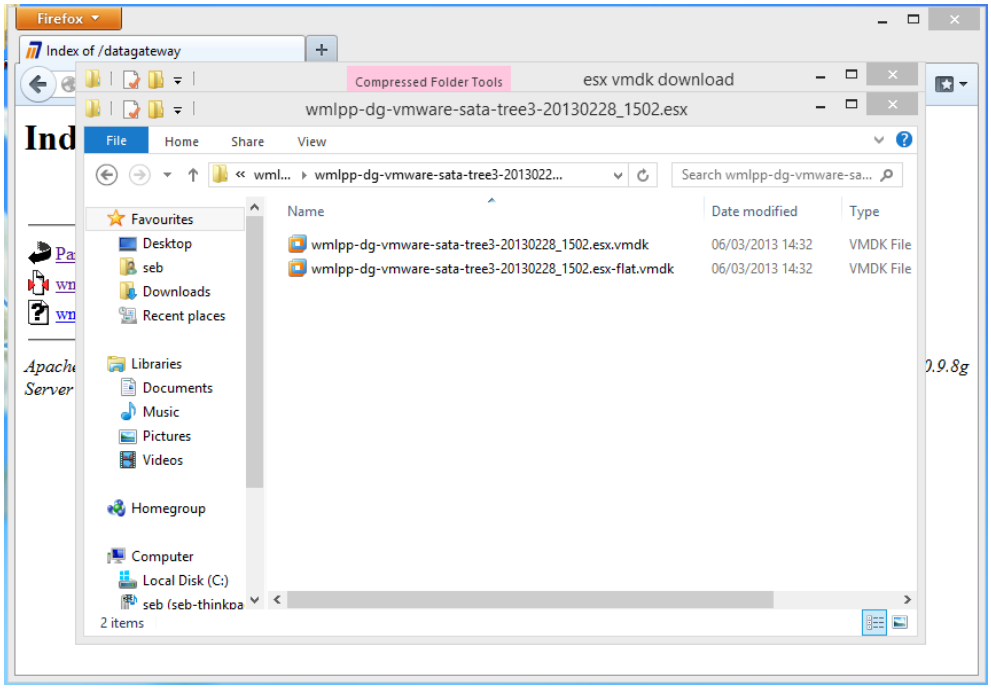

Upload both of the .vmdk files to the datastore of your ESXi host; in the **Summary** tab (Figure 2.41), right-click on the datastore and choose **Browse datastore**. The Datastore Browser has icons to upload two files. Your datastore will now contain the pre-installed virtual disk for your new virtual machine.

Figure 2.36. The "New Virtual Machine" wizard.

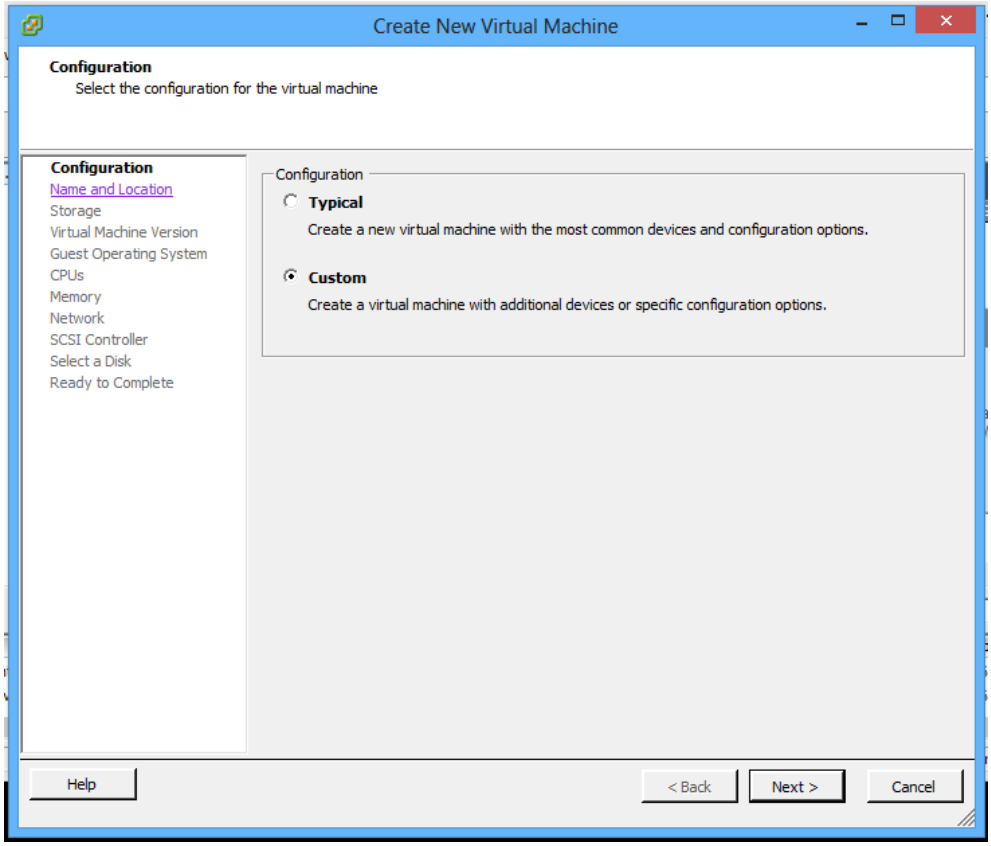

In the **Getting Started** tab, click the link **Create a new virtual machine**. The **Create New Virtual Machine** wizard will open (Figure 2.36). Select **Custom** so that it will be possible to specify that we want to use an existing, preformatted virtual disk. Click **Next**.

Figure 2.37. Choosing the name for the virtual machine.

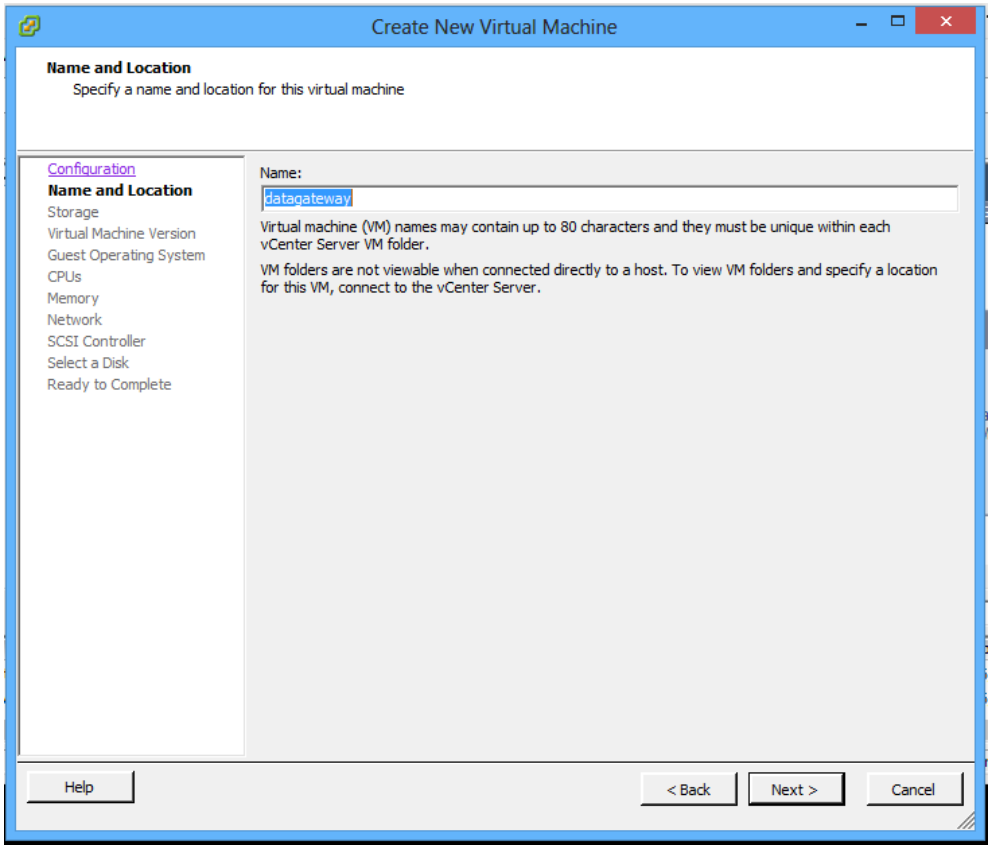

You can now choose the name for your virtual machine. Here, it has been named “datagateway” (Figure 2.37).

Figure 2.38. Choosing a datastore for the virtual machine.

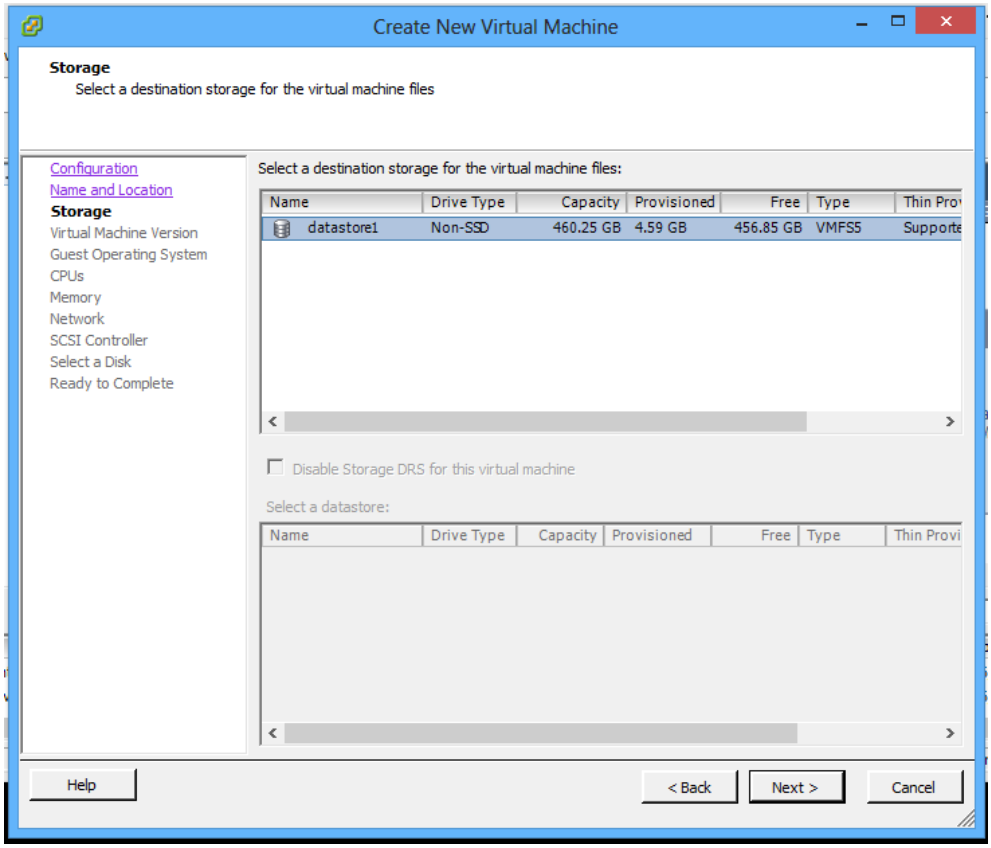

The next window allows you to specify a datastore in which you will create or find the virtual disk for the machine. Choose the datastore to which you uploaded the PdfPlus vmdk files. In this example, there is only a single datastore called datastore1.

Figure 2.39. Choosing the virtual machine hardware compatibility.

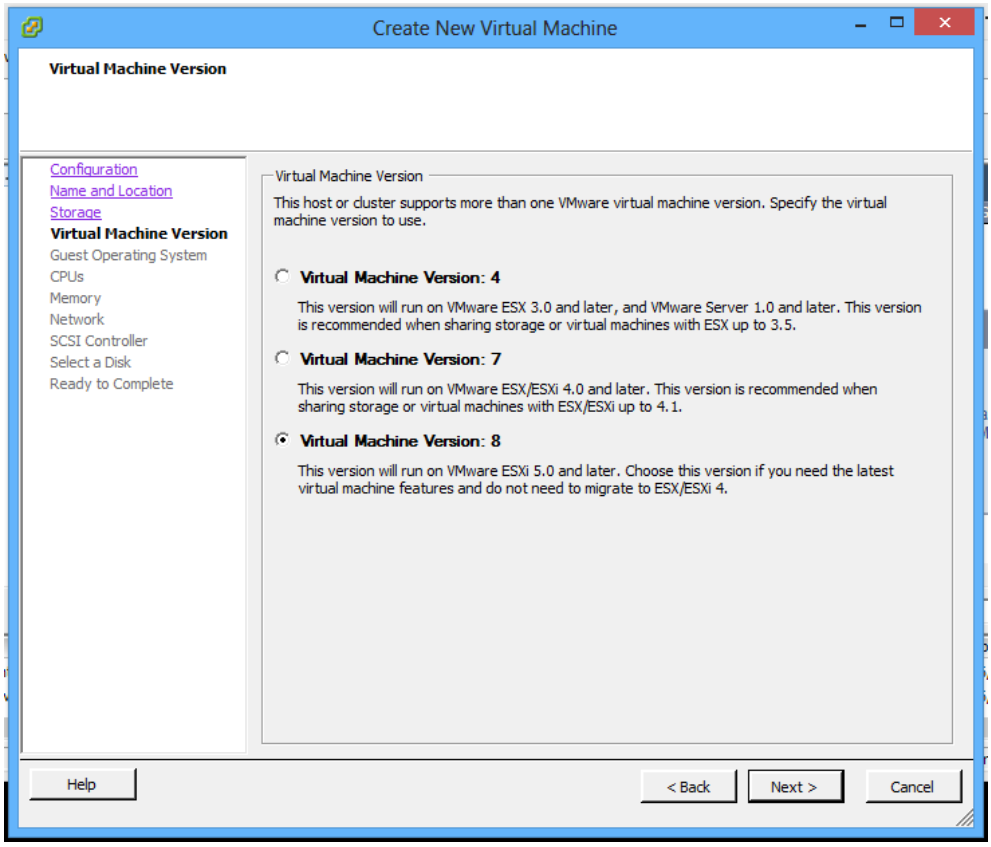

The Virtual Machine Version window allows you to specify the hardware compatibility for the virtual machine. This may be a consideration if you wish to move your virtual machine to older versions of VMware ESXi or to other VMware infrastructure products. In our example, we leave it at its default of **8**.

Figure 2.40. Specifying the operating system for the guest appliance.

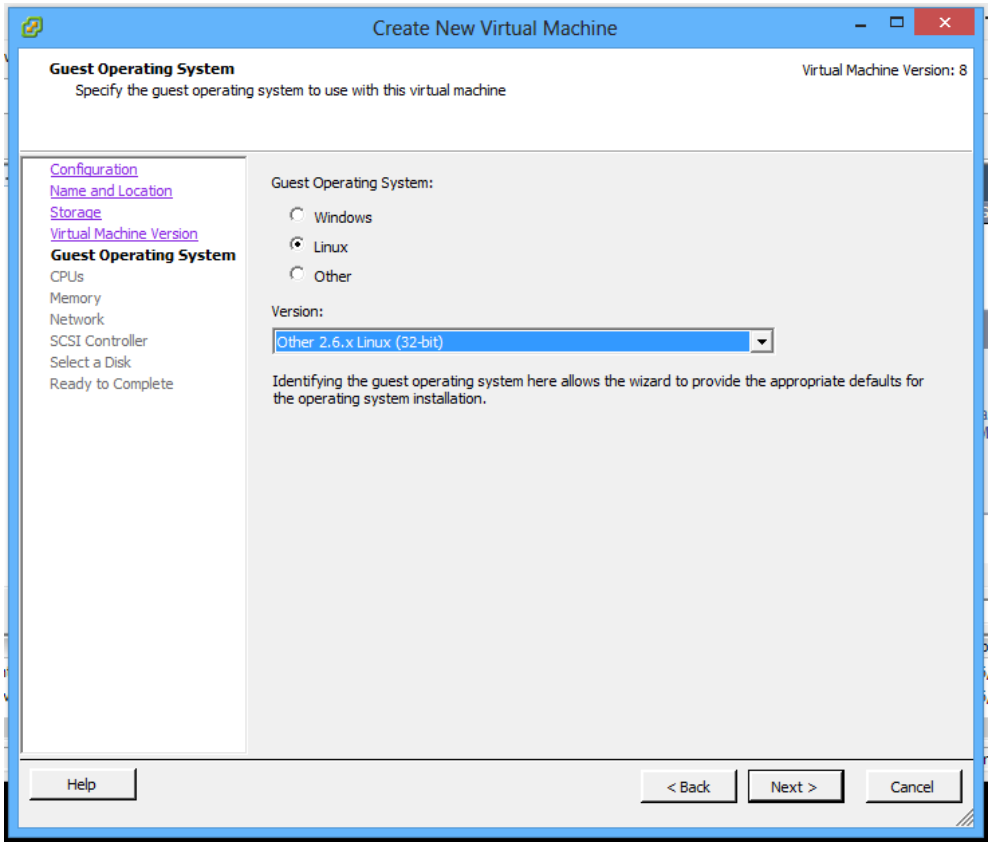

The next window (Figure 2.40) allows you to choose which operating system your target appliance will be running. All WML appliances run a 32 bit Linux operating system and so you should choose **Linux** for **Guest Operating System** and **Other 2.6.x Linux (32-bit)** for **Version**. Click **Next**.

Figure 2.41. Setting the number of processors for the virtual appliance.

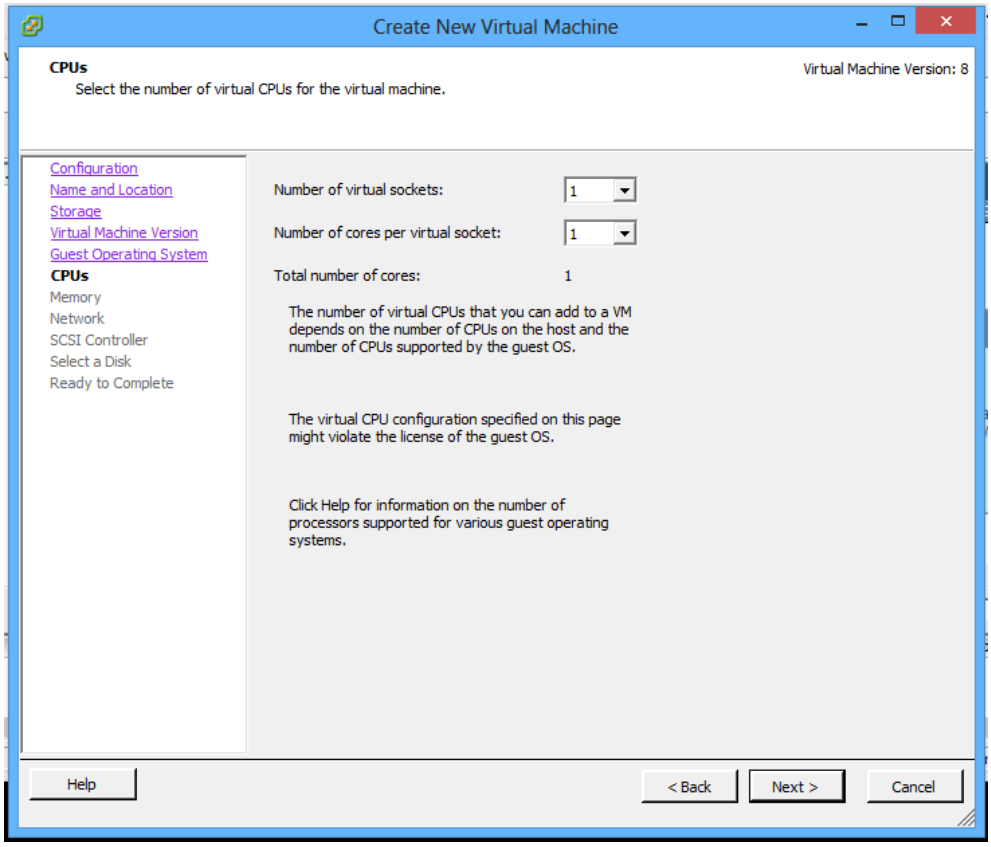

Figure 2.42. Setting the amount of RAM for the virtual appliance.

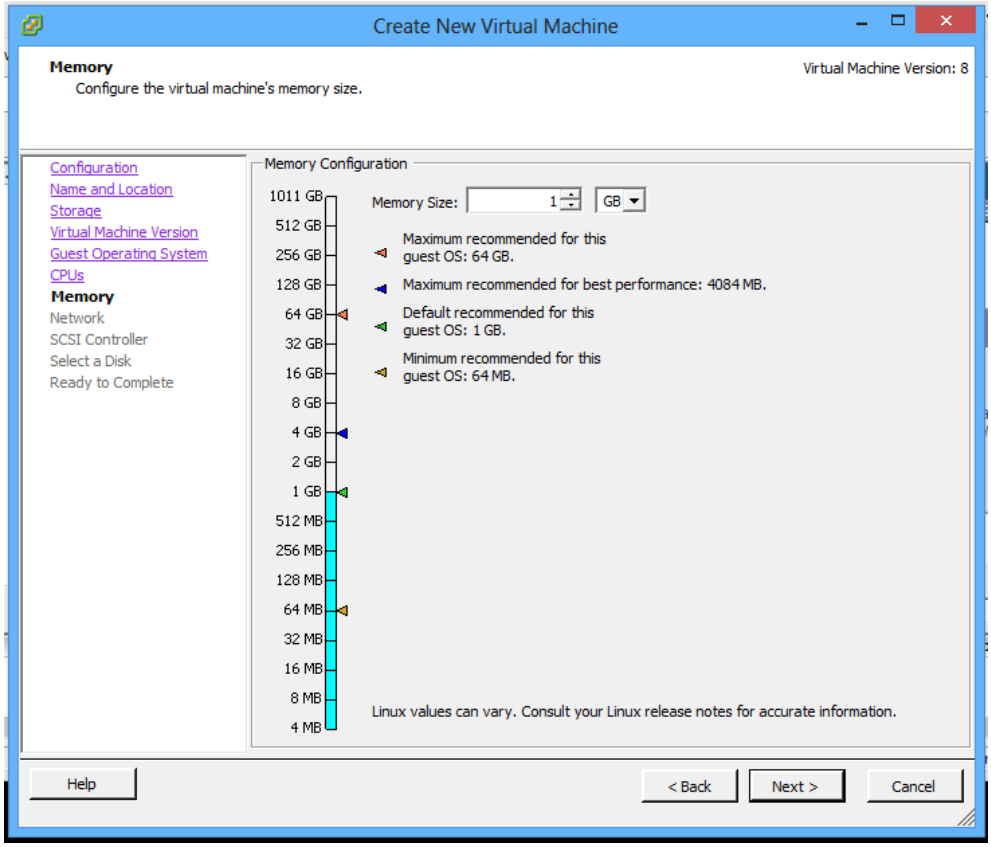

The next two screens (Figure 2.41 and Figure 2.42) allow you to specify the initial amount of RAM and processor resource to assign to your virtual print appliance.

Assign the desired number of processors for the virtual appliance. A single processor may be sufficient for simpler applications. Assigning two processors to the system will improve the responsiveness of the web interface when the print processing load is high.

Specify a minimum of 256 MB of RAM. Some applications may require more memory to run correctly, but 256 MB is sufficient for simpler tasks and short documents.

> **Note**
>
> The amount of RAM may be altered after the virtual appliance has been created. Altering the number of processors is not recommended by VMware, but it can be done. WML appliances will safely handle a change in the number of available processors.

Figure 2.43. The Network window.

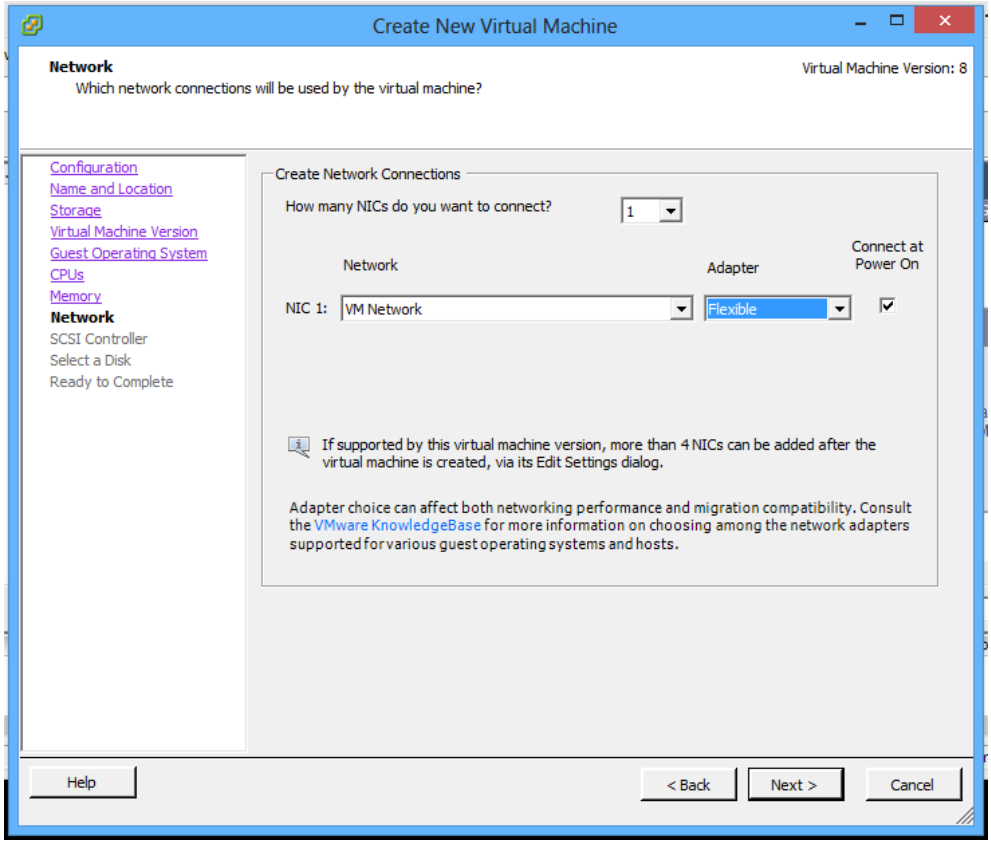

In the **Network** window (Figure 2.43) you can select the number of virtual network interface cards to assign to the virtual machine. The default of a single NIC is usually suitable for a PdfPlus. By default, the NIC for the virtual machine will join the same network as the ESXi host.

Figure 2.44. Selecting the virtual I/O controller type.

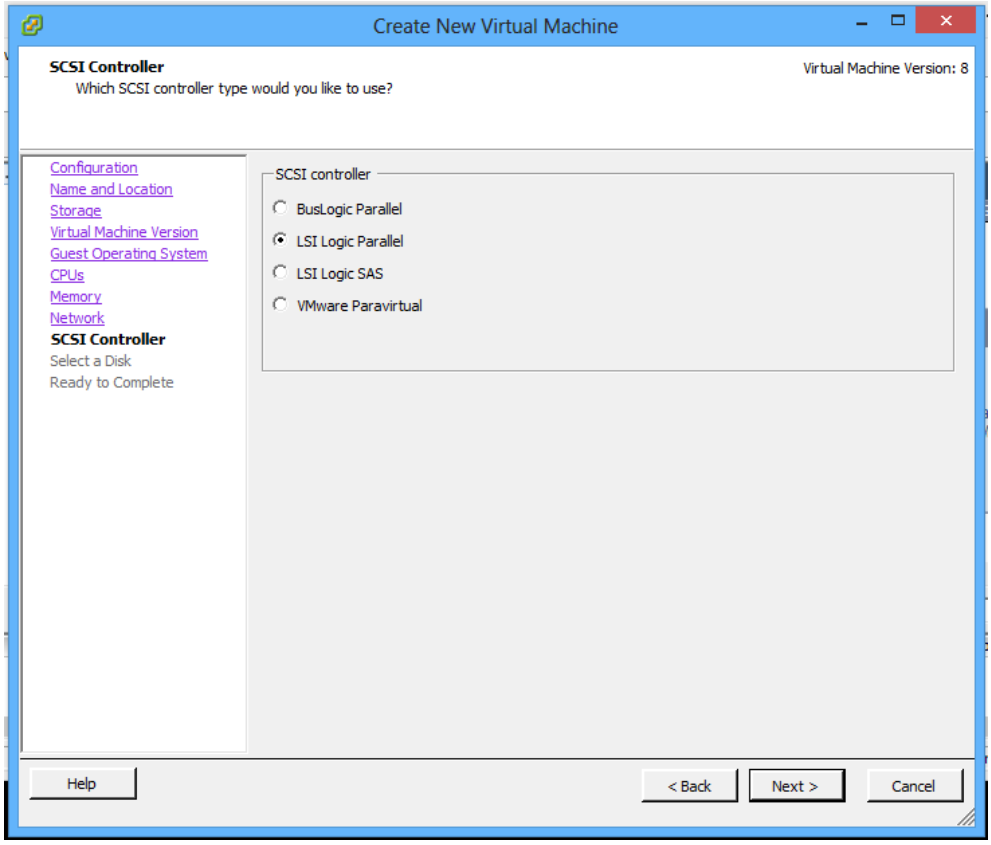

Figure 2.44 shows a screen allowing a choice of virtual SCSI disk controllers. Leave the default **LSI Logic Parallel** checked.

Figure 2.45. Selecting a hard disk for the virtual appliance.

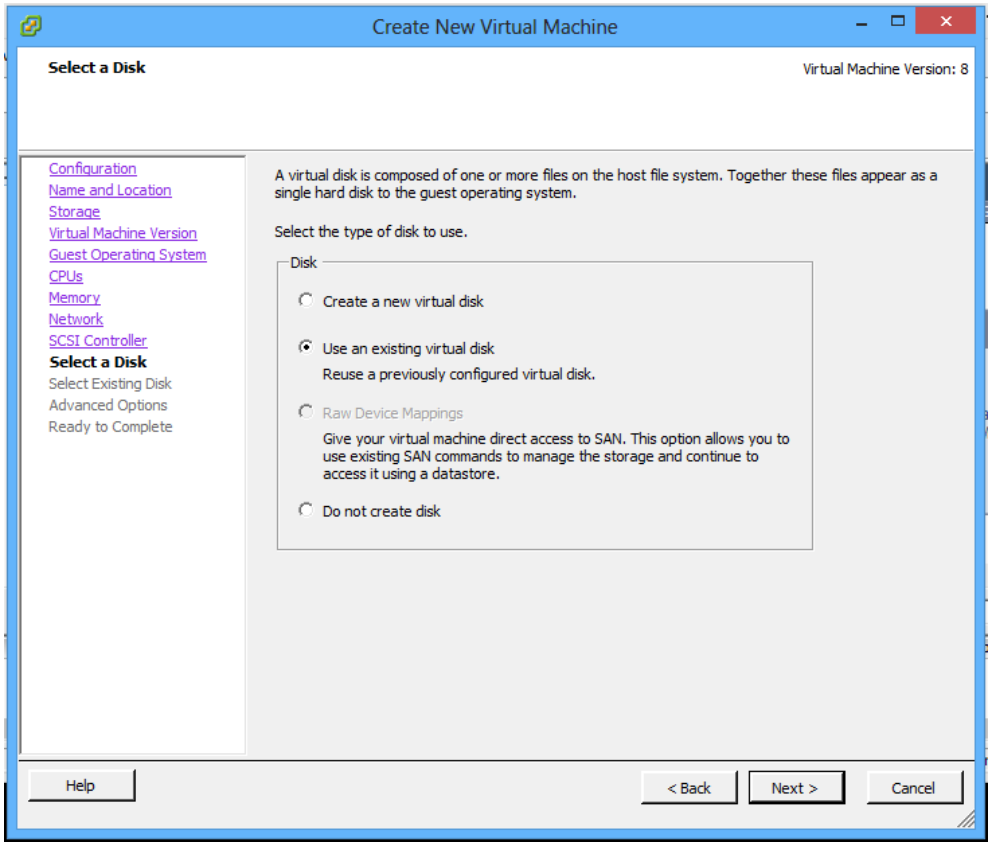

In the next window (Figure 2.45) you have the option to **Create a new virtual disk**. This would create an *empty* disk, which would then need to have the operating system installed (perhaps from a dvd or cdrom). Because you are using a pre-populated disk file supplied by WML, you should click on **Use an existing virtual disk**, which will open the next window, shown in Figure 2.46.

Figure 2.46. Selecting an existing disk.

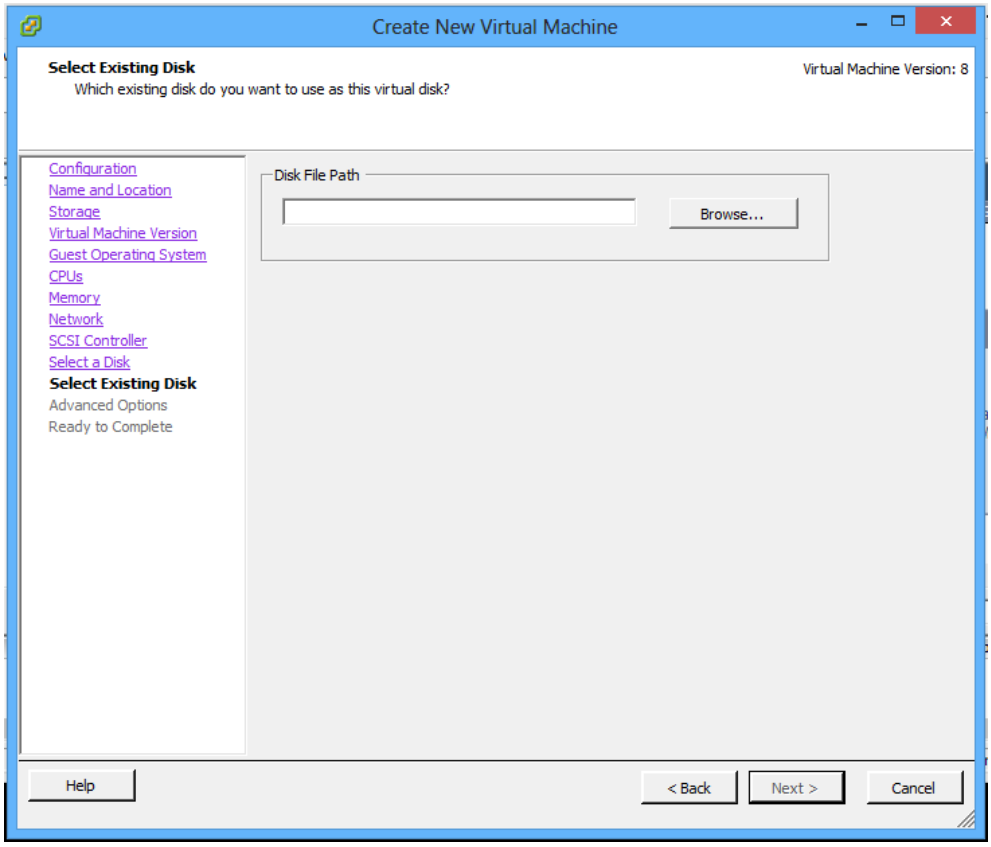

Figure 2.47. Selecting the WML Virtualized disk file.

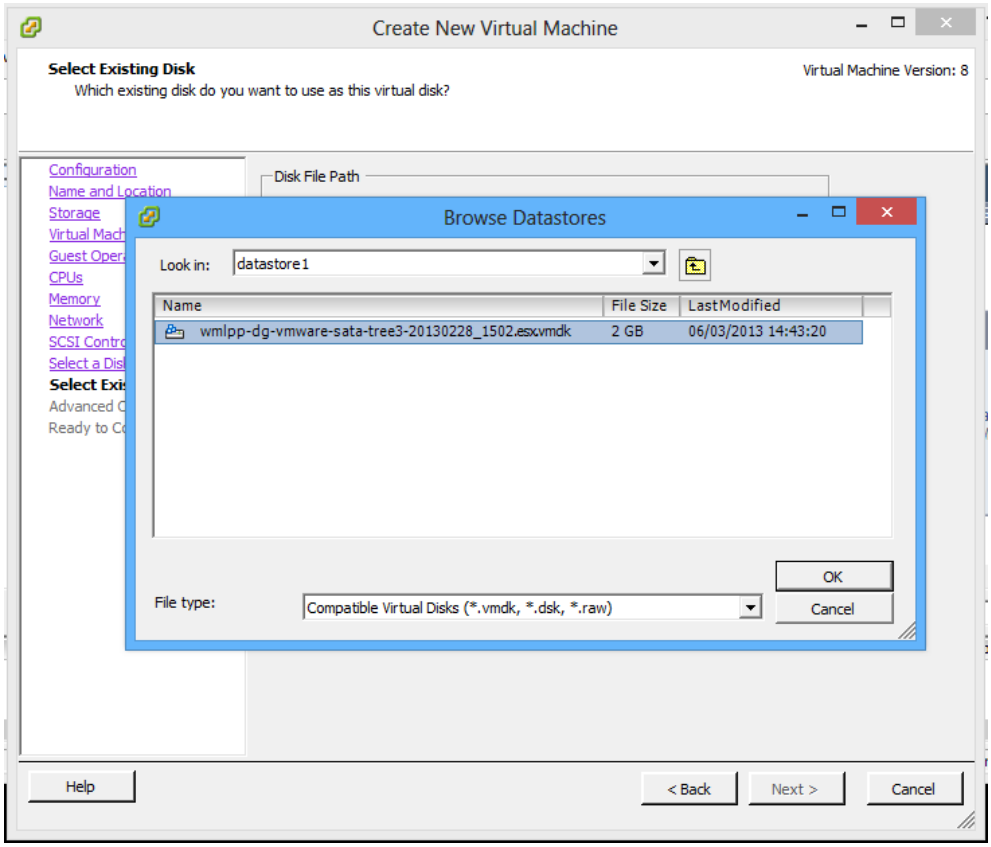

Browse to the vmdk disk image (Figure 2.47) in your datastore and select it.

Figure 2.48. Advanced disk options.

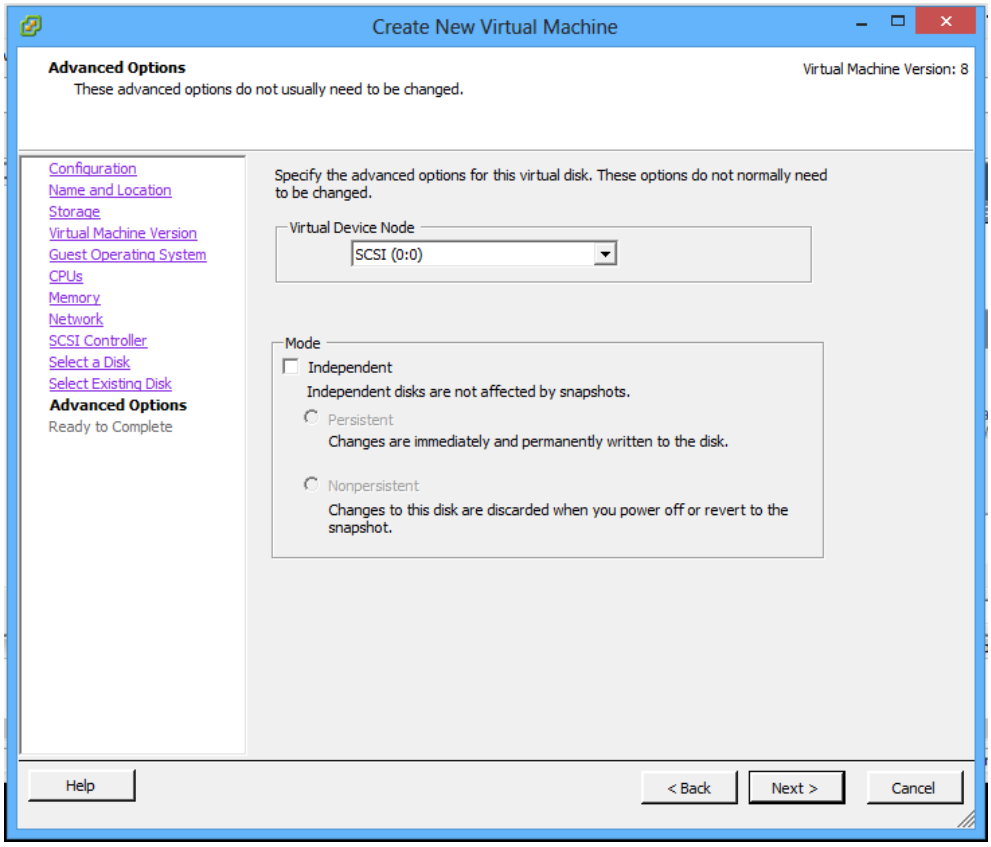

The penultimate window shows some advanced options for your virtual disk. Leave the virtual device node as **SCSI (0:0)** and click **Next**.

Figure 2.49. Ready to Complete - a summary of your new virtual appliance.

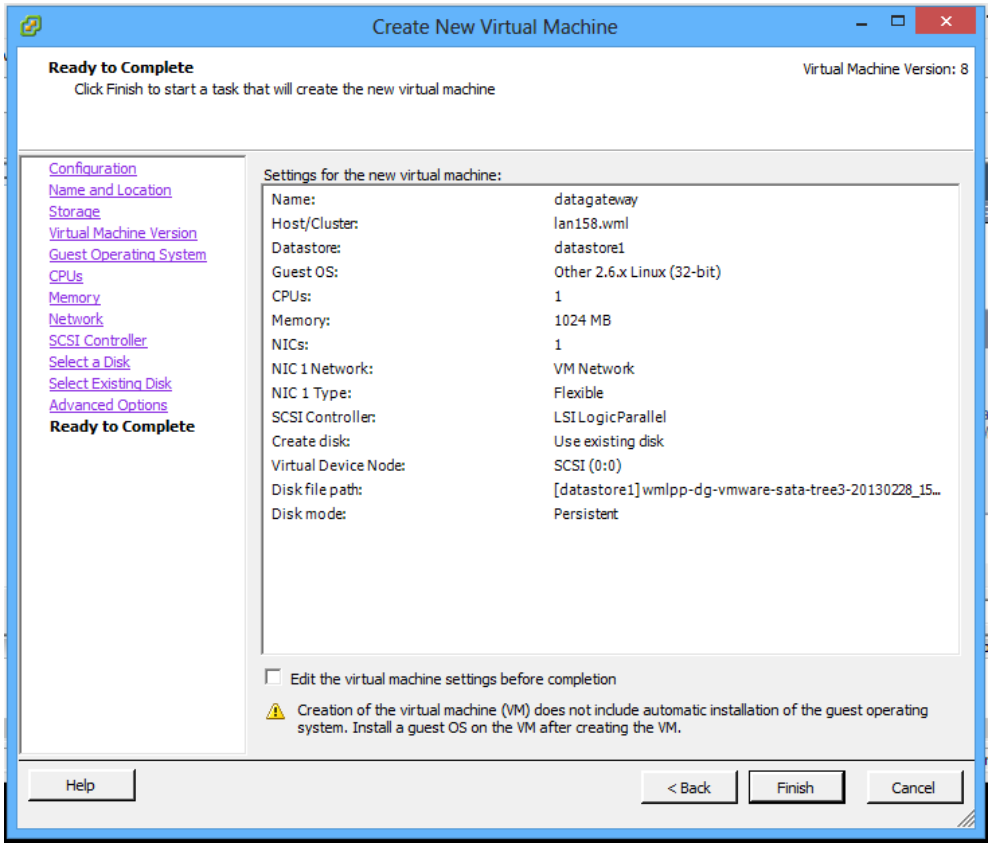

The last window shows a summary of the settings chosen for your new virtual machine (Figure 2.49). Click **Finish** to create the VM.

After creating your new virtual machine, a new item will appear in the inventory for your ESXi host (Figure 2.50). You can also access the console of your virtual machine in the vSphere client to find its IP address. Connect a browser to the IP address to access the Pdf-Plus web user interface.

> **Note**
>
> There is no need to install VMware tools on a WML Virtualized. VMware tools enhances the user experience when using a mouse within a graphical environment. The PdfPlus has no graphical console and hence VMware tools is not used.

Figure 2.50. The new appliance summary tab.

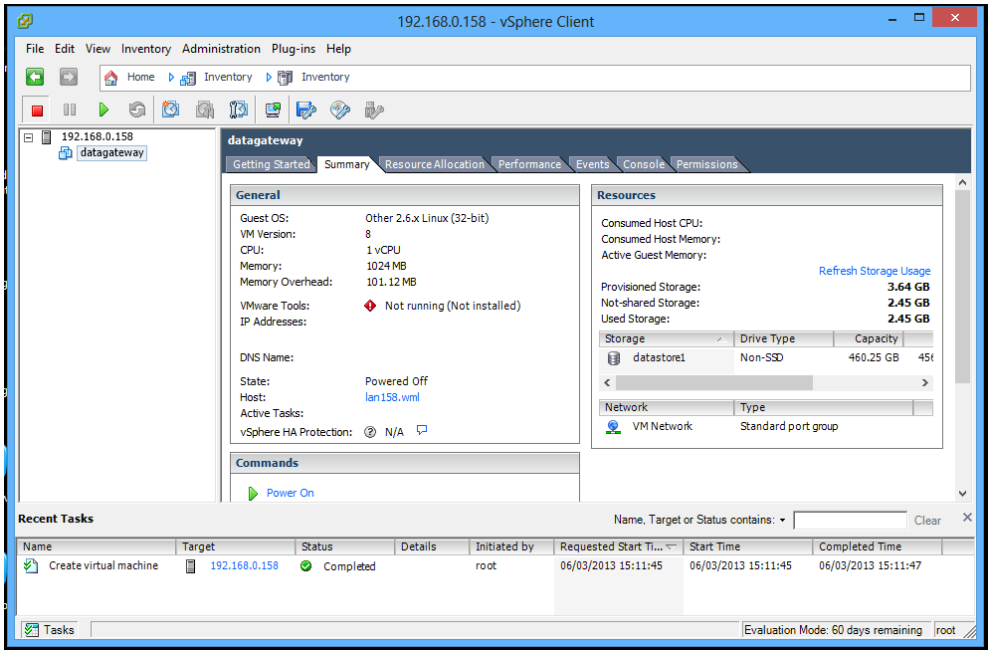

We hope you found this walk-through for VMware vSphere and ESXi useful.

2.2. Accessing the console

Occasionally, you may need to access the console of your PdfPlus, to carry out system configuration, maintenance and trouble-shooting.

Accessing the console on Tempest appliances

You will need a USB keyboard, a VGA monitor and an appropriate VGA cable to access the console of Tempest and Server-based appliances.

Connect a VGA monitor and USB keyboard to the PdfPlus. When the PdfPlus is switched on and the system has booted up you will see a login prompt.

You can log in as **admin** for manual IP address configuration or as **root** for other system tasks.

Accessing the console on Virtualized appliances

The virtual machine console for Virtualized appliances is available via the VMware Server management interface, the VMware Workstation UI or the vSphere client, depending on which VMware infrastructure is in use.

You can log in as **admin** for manual IP address configuration or as **root** for other system tasks.

Accessing the console on Server-based appliances

Please contact WML support for details of how to access the console on your Server-based PdfPlus.

2.3. IP address configuration

By default, the IP address of the PdfPlus will be allocated using DHCP.

To determine what IP address has been allocated you can consult your DHCP server logs. Alternatively, you can log into the PdfPlus console, view the IP address displayed there or carry out manual IP address configuration.

Console IP address configuration

When you need to configure the IP address without the web user interface, you can do so by logging in at the machine's console, as described in Section 2.2.

At the PdfPlus console, log in using the username **admin** and password **blankdoc** (or the password that has been set for the admin user on your system).

When you have logged in, the current network settings will be displayed. You will then be given the option to configure the IP address manually or via DHCP, or to quit.

```
Enter “manual” for manual IP configuration
or “dhcp” for automatic IP configuration (q to
quit).
```

If manual configuration is chosen, you will then be prompted for an IP address for the device, a Netmask address and, optionally, a gateway address.

```
Enter the IP for the device: 192.168.0.20
Enter the Netmask for the device: 255.255.255.0
Enter the Gateway (n for none): _
```

For both manual and DHCP configuration, you will be asked to confirm whether you want to reset the DNS settings. Specifying a DNS server IP which is not accessible on your network can lock the web based user interface of the PdfPlus.

```
Do you wish to reset the DNS settings? Specifying a DNS
server IP which is not accessible on your network
can lock the web based user interface of this device.
Enter 'y' to remove DNS settings, 'n' to keep them.
```

The PdfPlus will now be reconfigured, the new network settings will be displayed and the program will exit. To make further changes you will need to log in and follow the process through again.

2.4. Accessing the PdfPlus user interface

You can access the PdfPlus user interface from anywhere on your local network by directing a web browser at the IP address that has been allocated, either manually or via DHCP, to the device.

You will need to log in to administer and configure the PdfPlus.

PdfPlus users

The PdfPlus is configured with two levels of user: *System Manager* and *Print Manager*. The default passwords shipped with the system are “blankdoc” in each case.

The System Manager has unrestricted access to all configuration options on the PdfPlus.

The Print Manager has restricted access to the system, allowing the user to administer print jobs and print queues, but not make changes to system settings or operate the optional LAN Manager. For example, the Print Manager can't change the IP address of the PdfPlus or re-direct system logging messages.

Chapter 3. Disaster recovery and backup procedures

This chapter describes strategies for making a backup of your production PdfPlus so that a replacement system can be re-created in the event of it failing.

> **Warning**
>
> Please note that WML's standard maintenance contracts do NOT cover the re-build of a solution following a hardware failure unless virtual machine backup images can be provided or a WML disaster recovery system is in place.

3.1. Backup

Automated backup with WML's disaster recovery system

The best way to backup your PdfPlus is to implement WML's disaster recovery system. This is implemented as a feature in the Appliance Manager module (see Section 11.6).

The production system is configured as the "DR master" and has the Appliance Manager module installed. A second, backup PdfPlus is configured to act as the "DR backup". The DR master automatically pushes all configuration and archive files to the DR backup twice per day. It also makes a database dump of the mysql database and pushes this to the DR backup.

Crucially, every time a firmware upgrade is applied to the DR master, the upgrade is pushed to the DR backup, which automatically applies the upgrade itself. This means that the DR backup can be rebooted into production mode at any time, with the on-board configuration being no more than 12 hours out of date.

> **Warning**
>
> The DR backup system should be running on the same infrastructure as the main system. Either both should be KVM (Kernel Virtual Machine) based (such as those which are supplied by WML as hardware appliances) OR both should be VMware based. If one system is KVM based, and the second is VMware based, firmware upgrades applied to the main system will not propagate to the DR backup system.

Static virtual machine backup procedure

If a WML disaster recover system is not implemented, it is important to make backups of the virtual machine disk image files.

> **Note**
>
> All Server-based and Virtualized PdfPluss operate as virtual machines. Server-based systems are provided running on a pre-configured Ubuntu Linux server which provides a Linux KVM infrastructure for running the PdfPlus virtual machines.

Linux KVM virtual machines

If you have a Server-based PdfPlus or you have implemented PdfPlus on your own KVM infrastructure, you will be running a KVM based system.

If you are running PdfPlus on your own KVM infrastructure, please make regular backups of the PdfPlus qcow2 disk files, and of any files stored in a logical volume which is attached to the PdfPlus virtual machine.

> **Note**
>
> It is not necessary to shut down the PdfPlus while copying the qcow2 disk files, but it is advisable. If the PdfPlus is not shutdown when the backup is made, it will be necessary to run a file system check on the qcow2 file when the recovered PdfPlus boots up. Although this is completed automatically on every PdfPlus boot, there may be rare cases where the filesystem check cannot recover a file in the disk image.

If your KVM infrastructure has the facility to make complete backups or copies of running virtual machines, then consider making use of it.

If you are running a Server-based PdfPlus, you will need to backup the qcow2 disk files in **/var/vm/** on the Ubuntu server and any logical volume which has been presented to your PdfPlus as a second virtual disk.

To backup files in **/var/vm/**, use the scp (Secure CoPy) protocol to connect to the server on port 22 or 1522. Copy all files from the path **/var/vm/** on the remote system to your local backup location. You will need to have obtained and recorded the password associated with the user **wmldev** on the Ubuntu server.

> **Note**
>
> The WinSCP program (http://winscp.net/) is a free utility which implements the scp protocol and runs on Microsoft Windows.

You can determine if your PdfPlus is using an attached logical disk by consulting **System Settings > Status/Control > Share mounts**. This will open a dialog containing a table of **File Systems and Mounted Network Shares**. If you have a **Path** of **/dev/sdb** with **Mount Point /etc/persistent**, then you do have an attached logical volume.

Because this logical volume may be encrypted, the best way to back it up is to copy the files from the PdfPlus rather than via the underlying server.

Use scp to copy all files from **/etc/persistent/** *on the PdfPlus* to your local backup location. Use port 22 or 7522.

VMware virtual machines

If you are running your PdfPlus on a VMware infrastructure, please make use of the tools provided by VMware to make backups of your PdfPlus.

Database backups

An on-board MySQL database is used to record information for the Document Content Control module, the Electronic Document Delivery module and for full audit-trail reporting.

A static virtual machine backup will copy the existing state of the database, but it is also possible to make external backups of the database content.

To make dumps of the on-board database, ask your vendor to set up a read-only database user and then make use of the mysqldump utility to extract a snapshot of the database contents from the PdfPlus.

Making backups of PdfPlus system settings

The PdfPlus web user interface provides a method for saving configuration files, such as queue configuration and Document Content Control configuration.

These backups are particularly useful for creating duplicate systems for testing or to provide to WML to aid support. Each backup created is a tar archive file, containing the backed up files.

Refer to the section called “System settings” for details of backing up PdfPlus system settings.

> **Warning**
>
> This method does NOT allow you to backup archived documents, database content (used for Document Content Control, Electronic Document Delivery and audit report systems) or Archive Manager configuration.

3.2. Recovery

The method for recovering a PdfPlus following a failure of the hardware platform will depend on how the backups were made.

If the WML disaster recovery system is implemented, the DR master should be unplugged from the network, if it is still running, and the DR backup should be rebooted into “master mode”. (see Section 11.6 for more details).

If you are managing the virtualization infrastructure, then follow suitable steps to re-create a virtual machine from the backups you have made of the disk images/virtual machine.

If you are running a Server-based PdfPlus it will be necessary to have WML supply a pre-configured server running Ubuntu Linux with KVM onto which the virtual machine backups can be restored.

> **Note**
>
> The standard WML software maintenance contract will cover the work required to restore virtual machine disk image files onto the server. Hardware maintenance agreements should be arranged to cover the server hardware, which is otherwise supplied at additional cost.

If the system is a Tempest appliance, a replacement should be set up using the backup files made in the PdfPlus web user interface.

Chapter 4. PdfPlus Configuration

This chapter covers the steps necessary to configure the PdfPlus system. After configuration, the system will have the correct network settings, it will keep time correctly and system log messages will be logged to your syslog server.

After configuring the system, you will be ready to set up print queues, as described in Chapter 5: *Queue Management*.

This chapter also describes the configuration of passwords for authentication and describes how to modify the licence key.

4.1. System configuration

System configuration tools can be found in the PdfPlus web user interface by selecting the **System Settings** item from the main menu.

Network settings

To access the network settings controls (Figure 4.1), select **System Settings > Network**. This notation means "select the System button from the upper row, then select the Network tab in the lower row.".

Figure 4.1. Network settings

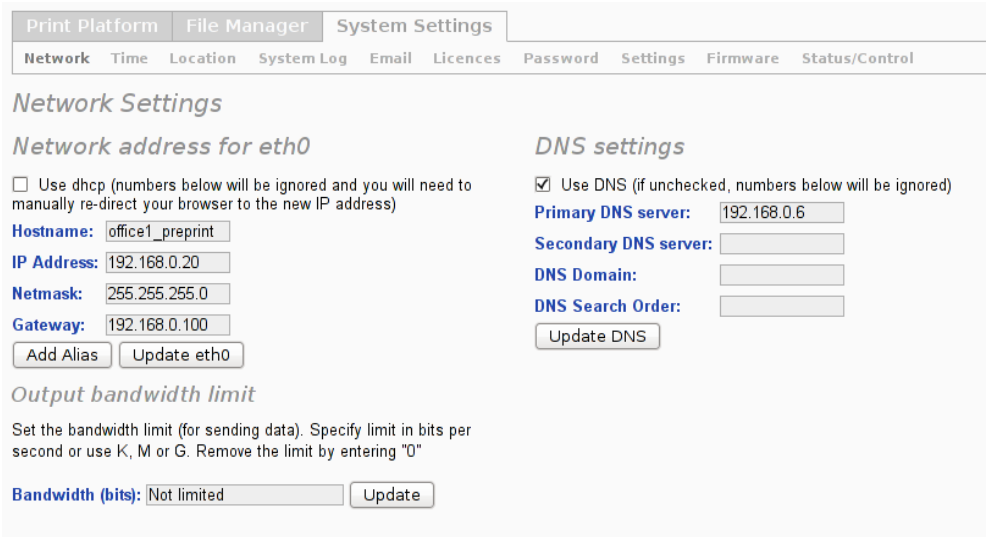

Here you can set the network address for the network interface(s). Figure 4.1 shows a PdfPlus on which there is one network interface, **eth0**. On hardware platforms with multiple

network interfaces, additional sections for the extra interfaces will be shown. The interfaces are named **eth0**, **eth1**, **eth2** and so on.

The parameters for each network interface can be set to fixed, static values or can be set to be dynamically updated using DHCP (Dynamic Host Configuration Protocol). If you check "Use dhcp", the IP address will be allocated by your DHCP server. The client ID passed to the DHCP server will be "eth0" for the **eth0** interface, "eth1" for the **eth1** interface and so on. You may need to consult the log of your DHCP server to find the IP address allocated to the PdfPlus. Alternatively, access the console of the PdfPlus and login as admin, as described in Section 2.2 in Chapter 2: *Getting Started*. On doing so, the current network settings will be displayed.

Configuring multiple aliases

The PdfPlus can be configured with multiple IP addresses for each network interface. To add an alias, press the "Add Alias" button. A field will appear in which you can specify a static IP address or enter **dhcp** to have the alias obtain a dynamic IP address from your DHCP server.

When using DHCP to allocate IP addresses to multiple aliases on the same network interface, you will need to configure the server to make use of the *client ID* parameter. The client ID is an identifier passed by the DHCP client to the DHCP server to identify the request. The first alias on the interface **eth0** has the client ID **eth0:1**, the second alias has the client ID **eth0:2** and so on.

Configuring DNS settings

DNS settings are also configured on the Network Settings page. You can configure the primary and secondary DNS servers, along with the domain and a domain search order for name resolution. These parameters are placed in a standard Unix/Linux **resolv.conf** file.

Bandwidth limit

It is possible to limit the rate at which the PdfPlus sends data on its output. This feature can be used to preserve a portion of a site's bandwidth for other uses which may be sensitive to latency.

> **Note**
>
> It is not possible to limit the rate at which data comes *into* the appliance. This is a fundamental limitation of the IP protcol, which is designed to accept all incoming packets on an interface. Bandwidth limiting is always achieved by limiting the rate at which data packets are sent.

To limit the bandwidth for a given interface, enter a number in the corresponding "Bandwidth (bits)" field and press "Update". A suffix of K for kilobits per second, M for megabits per second or G for gigabits per second can be appended.

To remove the bandwidth limit, enter 0 in the "Bandwidth (bits)" field.

System date and time

To access the system date and time controls (Figure 4.2), select **System Settings > Time.**

Figure 4.2. Setting system date and time

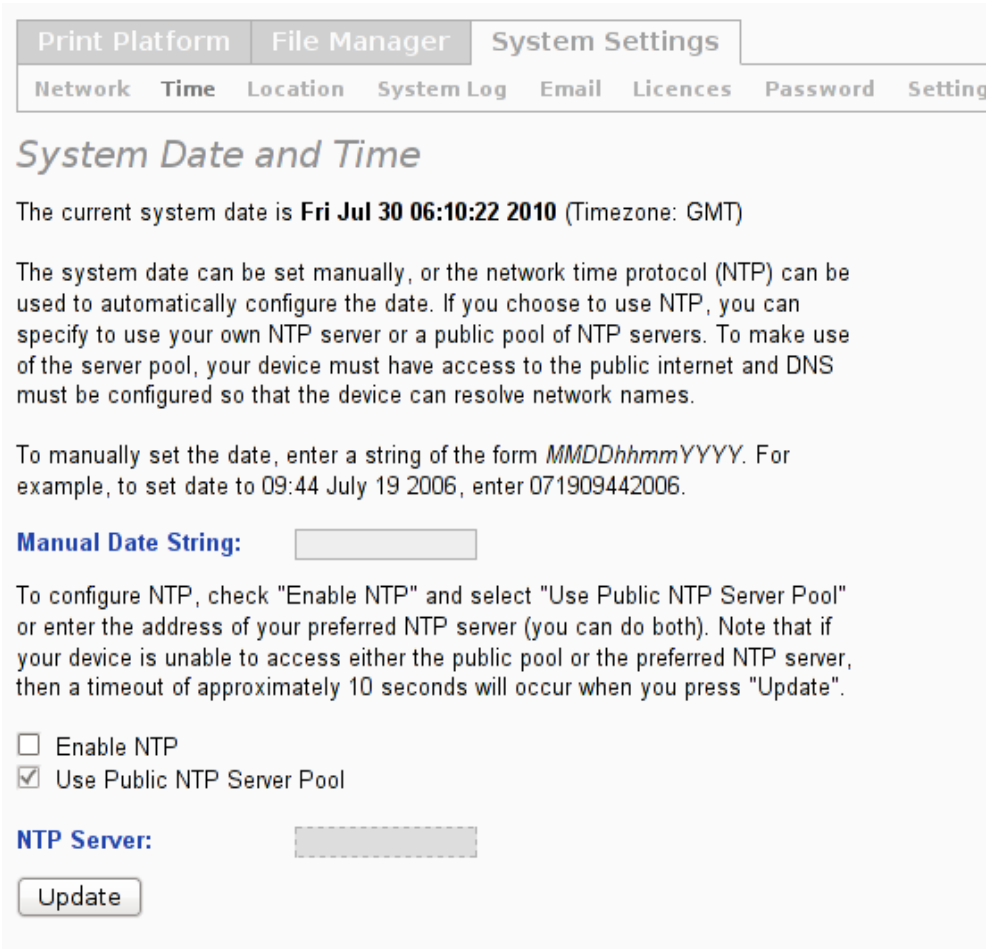

The system date can be set manually or using network time protocol (NTP). If you choose to use NTP, you can specify to use your own NTP server or a public pool of NTP servers. To make use of the server pool, your device must have access to the public internet and DNS must be configured so that the device can resolve network names.

To set the date, manually enter a string of the form **`MMDDhhmmYYYY`**. For example, to set date to 09:44 July 19 2006, enter **`071909442006`**.

To save your changes, click on the “Update” button.

To configure NTP, check “Enable NTP” and select “Use Public NTP Server Pool” or enter the address of your preferred NTP server (you can do both). Note that if your device is unable to

access either the public pool or the preferred NTP server, then a timeout of approximately 10 seconds will occur when you press “Update”.

Don't forget to set the correct timezone (Section : System timezone and locale), to ensure that daylight savings adjustments are automatically applied.

System timezone and locale

To access the system timezone and locale controls (Figure 4.3), select **System Settings > Location**.

Figure 4.3. Setting system timezone and locale

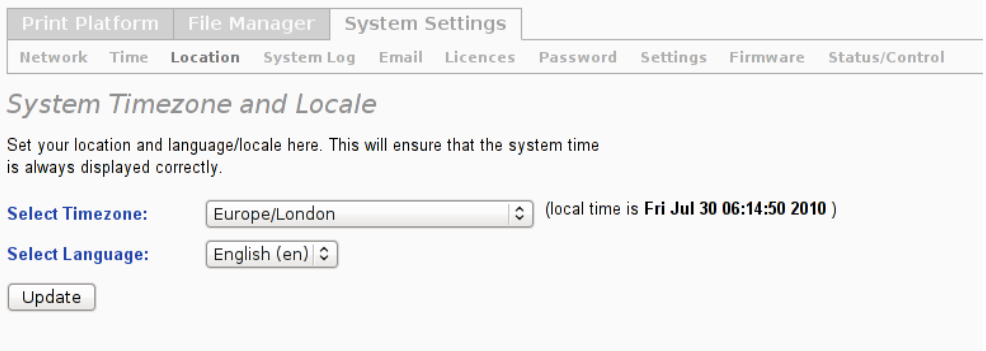

The drop-down controls can be used to select your timezone and language/locale.

It is important for time-keeping (see “System date and time”) to set the the correct time-zone in which you are operating.

To save your changes, click on the “Update” button.

System logging

To access the system logging controls (Figure 4.4), select **System Settings > System Log**.

The system logging controls allow you to choose whether or not to permanently store system log messages. As well as being useful for diagnosing problems system reports are

If left unconfigured, the PdfPlus will only store system messages to a temporary log in RAM. This log is frequently truncated and overwritten. In order to permanently store the system log, messages may be sent to a syslog server. Messages may also be stored to a network share or local drive (on Server-based and Virtualized systems).

Figure 4.4. System logging

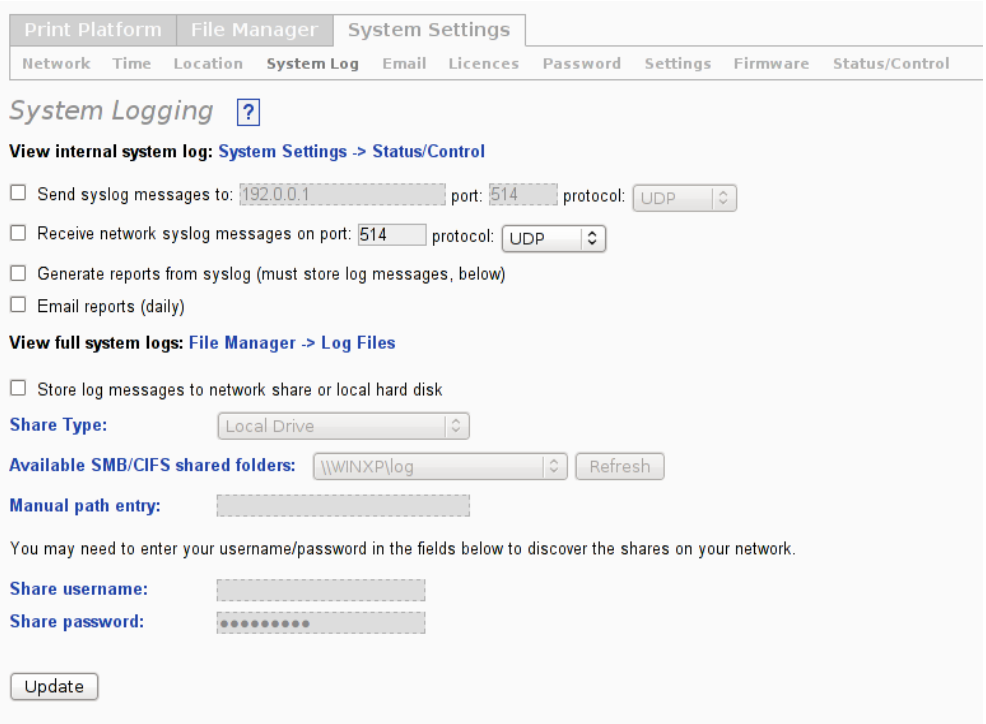

If you want to send syslog messages to another device, check the "Send syslog messages..." box and enter the IP address, port number and protocol to be used. You may send the messages to multiple devices by entering multiple, comma-separated IP addresses.

You may also choose to receive network syslog messages on the PdfPlus by checking the "Receive network syslog messages..." box and filling in the port number and protocol appropriately.

To store messages to a network share or local disk, check the "Store log messages to network share..." box and specify the storage type as described in the section called "Selecting network/hard disk storage".

> **Note**
>
> If you un-check "Send...", "Receive..." and "Store...", system log messages will still be written to a file in RAM, but they will be periodically truncated and will not persist after a reboot.

Permanently stored system log messages are stored to two files, wmlpp_job.log and wmlpp_sys.log. As the file names suggest, wmlpp_job.log contains messages relating to print jobs, while wmlpp_sys.log stores messages relating to other aspects of the system. Example 4.1 shows an example set of log files.

Example 4.1. Example system log files

```
2009-07-10-wmlpp_job.log          2009-07-12-wmlpp_sys.log
2009-07-10-wmlpp_job.log.partner  2009-07-13-wmlpp_job.log
2009-07-10-wmlpp_report.csv       2009-07-13-wmlpp_report.csv
2009-07-10-wmlpp_sys.log          2009-07-13-wmlpp_sys.log
2009-07-11-wmlpp_job.log          2009-07-14-wmlpp_job.log
2009-07-11-wmlpp_job.log.partner  2009-07-14-wmlpp_job.log.partner
2009-07-11-wmlpp_report.csv       2009-07-14-wmlpp_sys.log
2009-07-11-wmlpp_sys.log          wmlpp_job.log
2009-07-12-wmlpp_job.log          wmlpp_report.csv
2009-07-12-wmlpp_job.log.partner  wmlpp_sys.log
2009-07-12-wmlpp_report.csv
```

wmlpp_job.log and wmlpp_sys.log are "rotated" daily in the early morning. Each day, the files are renamed in the format YYYY-MM-DD-wmlpp_job.log (for wmlpp_job.log) and the current log is then truncated. The system messages for the *previous* day are stored in the dated file. For example, 2009-07-12-wmlpp_job.log contains the job logs for 2009-07-11.

To save your changes to the system logging settings, click on the "Update" button.

Selecting network/hard disk storage

The "Network/Hard Disk Storage" section enables you to specify a network share or local drive (on Server-based systems) to use for system log storage. The network share may be either a Unix/Linux Network File Server (NFS) share or a Windows/Samba (SMB/CIFS) network share. The preferred storage option may be selected using the "Share Type" drop-down.

Selecting a share

NFS (Linux/Unix shares)

To connect to a Linux/Unix NFS share, choose **NFS (Linux/Unix)** from the **Share Type** drop-down. Enter the path to the NFS share in the following format:

`[Network/IP Address]:[folder path]`

For example:

`192.168.100.100:/input_data`

The NFS server's exports file should be configured to allow both read and write access to the directory path (/input_data in the example). An example of a suitable line in /etc/exports on the NFS server is:

```
/input_data 192.168.0.0/255.255.255.0(rw,no_root_squash,sync,subtree_check)
```

If the NFS share is to be used as an *archive* Share, then correct permissions must be set on the share directory on the NFS server. The following would be required for an archive share specified as:

```
192.168.100.100:/archiveData
```

On the server it would be necessary to change the user on this directory to "root" or 0 and the group of the directory to 7, which is the GID for the "lp" user on some systems:

```
root@nfsServer# chown 0:7 /archiveData
```

Finally, ensure that both user and group are able to read and write to the directory:

```
root@nfsServer# chmod ug+rwx /archiveData
```

SMB/CIFS (Windows shares)

For Windows SMB/CIFS network shares, select **SMB/CIFS (Windows share)** and click on the **Refresh** button to scan the network for available shares. This will populate the **Available SMB/CIFS shared folders** drop-down.

> **Important**
>
> *Most recent Microsoft Windows file servers will not broadcast their available shares to the network without authentication.*
>
> If necessary, enter your network username and password in the fields **Network user name** and **Network password** and then press **Refresh**. The username and password will be used to obtain the list of available shares.

It is also possible to enter the share path manually in the **Manual path entry** text field, by choosing the option **Enter path in text box** from the **Available SMB/CIFS shared folders** drop-down.

Enter the SMB/CIFS path in the format:

```
\\MACHINE\shareName
```

The **Network user name** and the **Network password** fields must be filled in to attach shares which require authentication.

To save your changes, click on the "Update" button.

Email settings

To access the email settings controls (Figure 4.5), select **System Settings > Email**.

Figure 4.5. Email settings

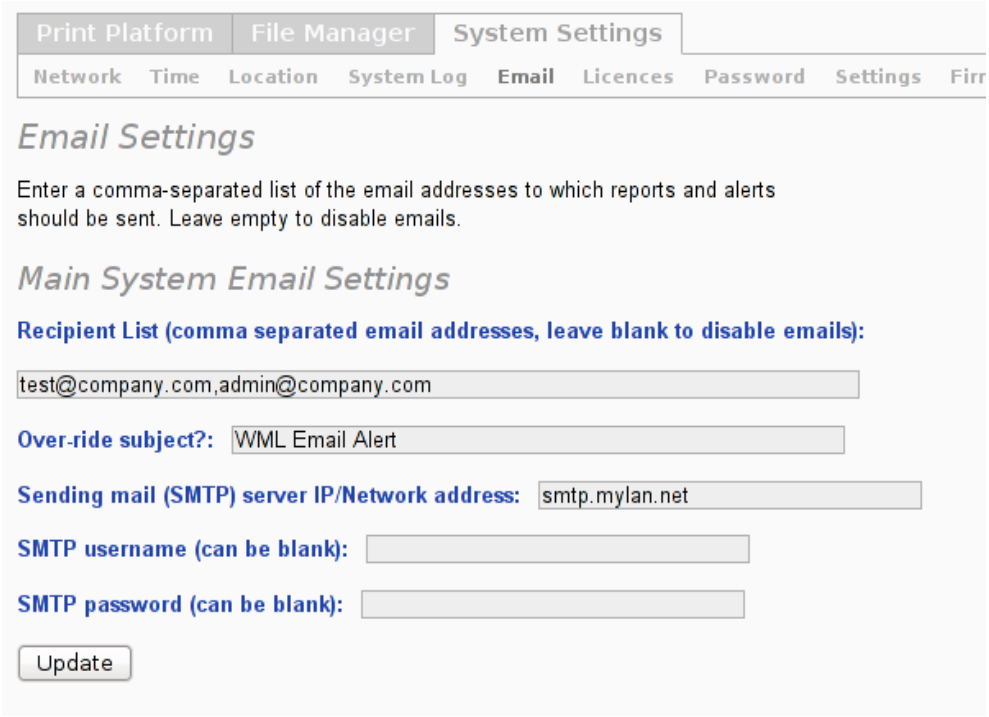

The system may be set up to send reports and alerts to specified email addresses. Enter a comma-separated list of email addresses, or leave the field blank if you do not wish to receive system emails.

You will also need to provide the IP/network address of the SMTP server to be used to send the emails and SMTP username and password, if required.

To save your changes, click on the "Update" button.

Licence key

To access the licence key controls (Figure 4.6), select **System Settings > Licences**.

Figure 4.6. Setting the licence key

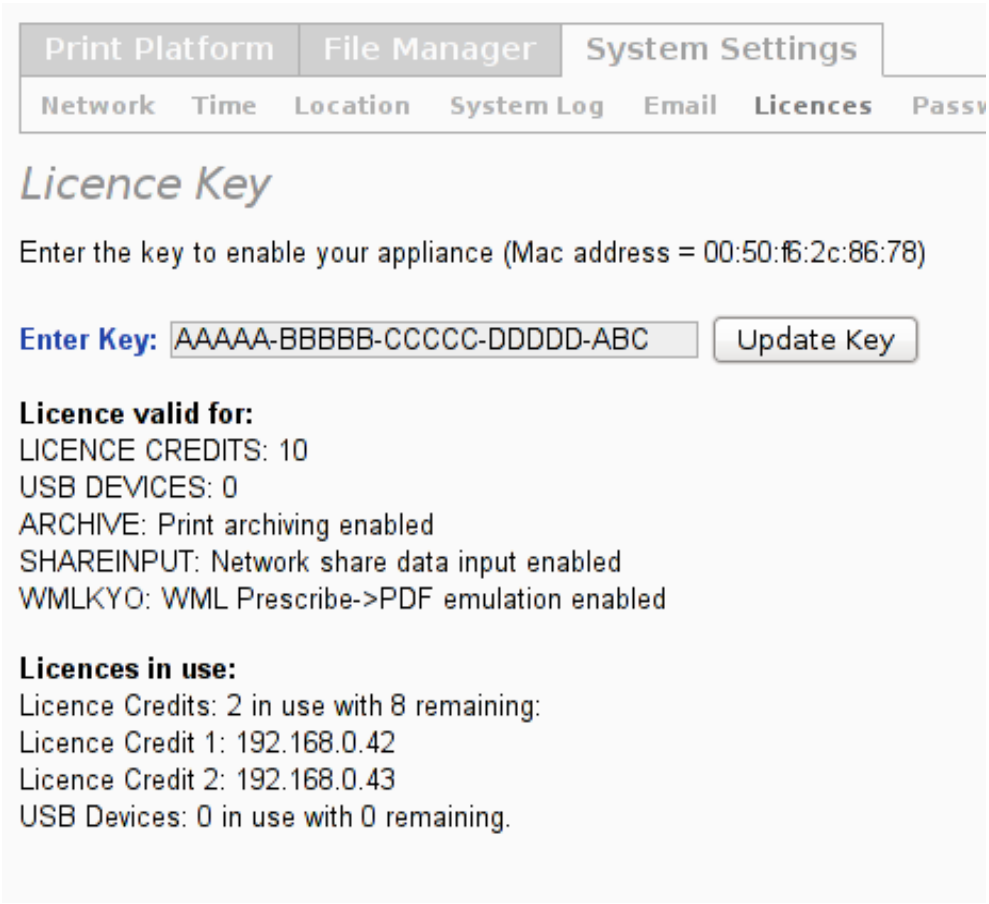

Enter the licence key received from the appliance vendor and click on "Update Key". The licence details will be displayed below.

WML Print Platform Licences have a number of "licence credits" enabled. They also provide a number of flags to enable certain features on the appliance, such as the Prescribe emulation and archiving. A number of USB devices may also be enabled in the licence.

A licence credit enables the use of an IP address on the appliance or the use of a input share/hot folder or archive share (if enabled with the relevant flags). The number of valid licence credits, USB devices and enabled feature flags is shown for the key which is stored on the appliance. The number of licence credits and USB devices which are in use will also be displayed.

System passwords

To access the system password controls (Figure 4.7), select **System Settings > Password**.

Figure 4.7. Setting system and print manager passwords

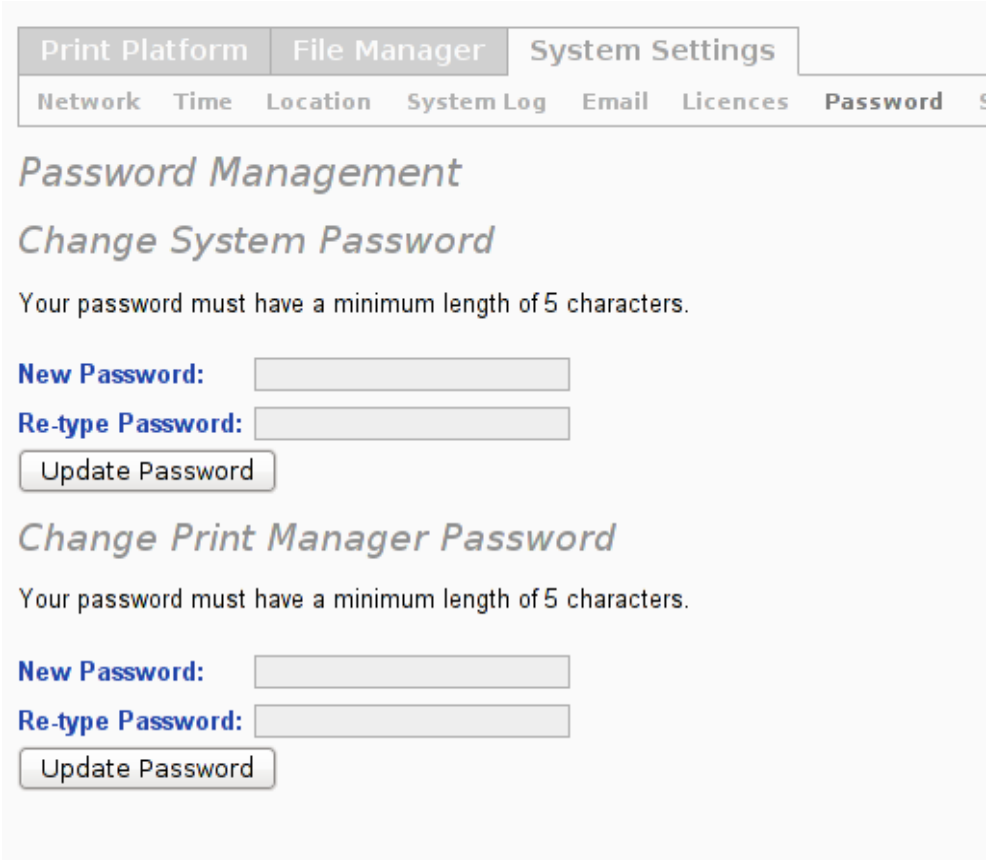

You can use these controls to set passwords for the system and print manager users.

To change either password (you can't change both at once), enter the new password in both boxes and click on “Update Password”.

> **Note**
>
> Passwords must have a minimum length of five characters.
>
> The default password is **`blankdoc`**.

System settings

To access the controls to backup, restore and reset system settings (Figure 4.8), select **System Settings > Backup**.

Figure 4.8. System settings

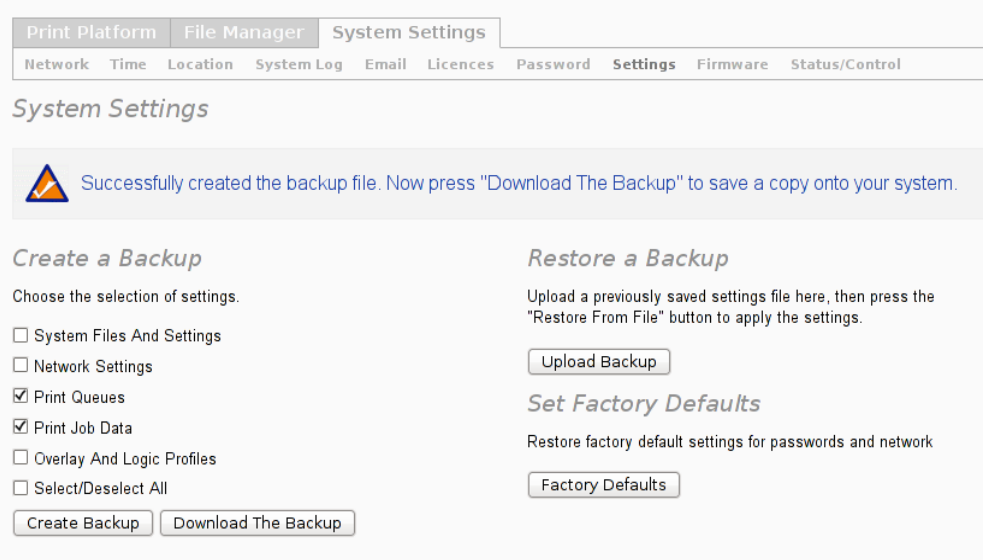

Backing up system settings

Before backing up the system settings decide which parts of the configuration need to be stored. The options for back up will include:

- System Files and Settings
- Network Settings
- Print Queues
- Print Job Data
- Overlay and Logic Profiles

There is also a checkbox to select or deselect all options.

It is common to want to copy the print settings from one appliance to another, without making any other changes to the target system. In this case, a backup of “Print Queues” and “Overlay and Logic Profiles” would be a suitable selection.

To copy an entire system setup, but without affecting the network settings of the target appliance, select all options *except* “Network Settings”.

To make a full backup to use in the case that the hardware failed and an identical setup were required on new hardware, select all options.

After making the selection, click “Create Backup”. A new button “Download Backup” will appear. Click this button and save the file to the local system.

Click on the link provided to backup the system settings. You will be prompted to save a file called “settings.tar”. This contains an archive of all the current settings for the device.

Restoring system settings

To restore settings from a previously saved file, click “Upload Backup”, and browse to the file. After the file uploads successfully, a new button “Restore From File” will appear. Click this button to apply the uploaded backup.

Restoring factory defaults

To restore factory default settings for passwords and network, click on “Factory Defaults”. A confirmation dialog box will appear.

Firmware

To access the system firmware controls (Figure 4.9), select **System Settings > Firmware**.

Figure 4.9. Firmware

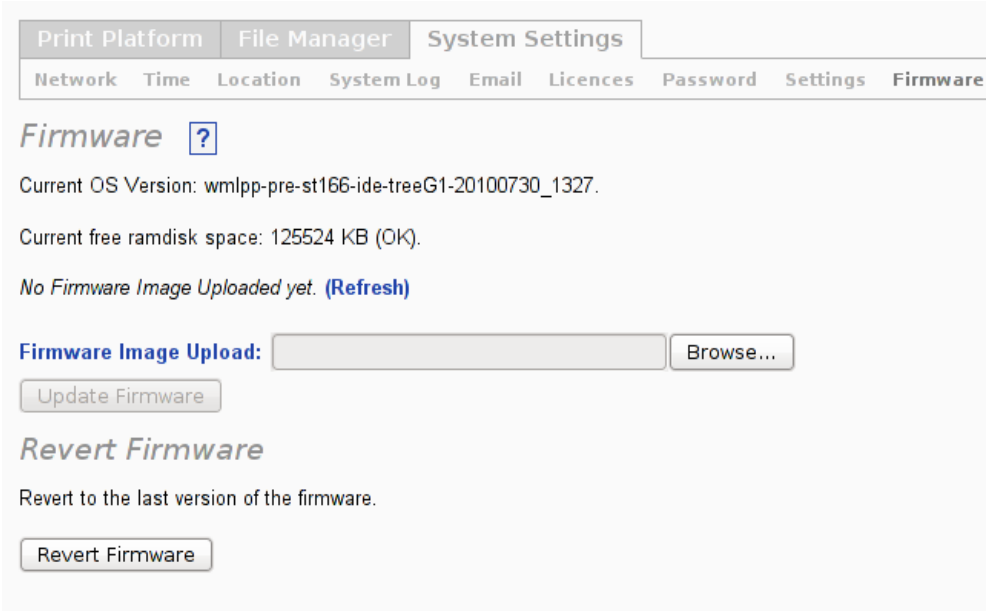

To perform a firmware upgrade, you must first upload a new firmware image to the device, either by browsing to the file, by uploading a firmware image into the “update” directory via FTP or by using scp/sftp to copy the firmware image into the “/tmp/uploads/update” directory. Details of the current and uploaded firmware images are shown on the page. Click on “Refresh” if you have uploaded an image and it is not yet displayed.

When the image has been uploaded, the “Update Firmware” button will be displayed. Click on the button to start the update process. The process will reboot the appliance. Avoid interrupting the firmware update, which takes approximately 45 seconds in total. The brower will re-connect to the user interface when the update has completed.

To revert to the previous version of the firmware, click on the “Revert Firmware” button.

Status and Control

The system status and control page provides status information such as system load and memory usage, the ability to shut down or reboot the appliance and a feature to allow WML staff to access the appliance for support.

To access the system status and control page (Figure 4.10), select **System Settings > Status/Control**.

Figure 4.10. Status and Control user interface

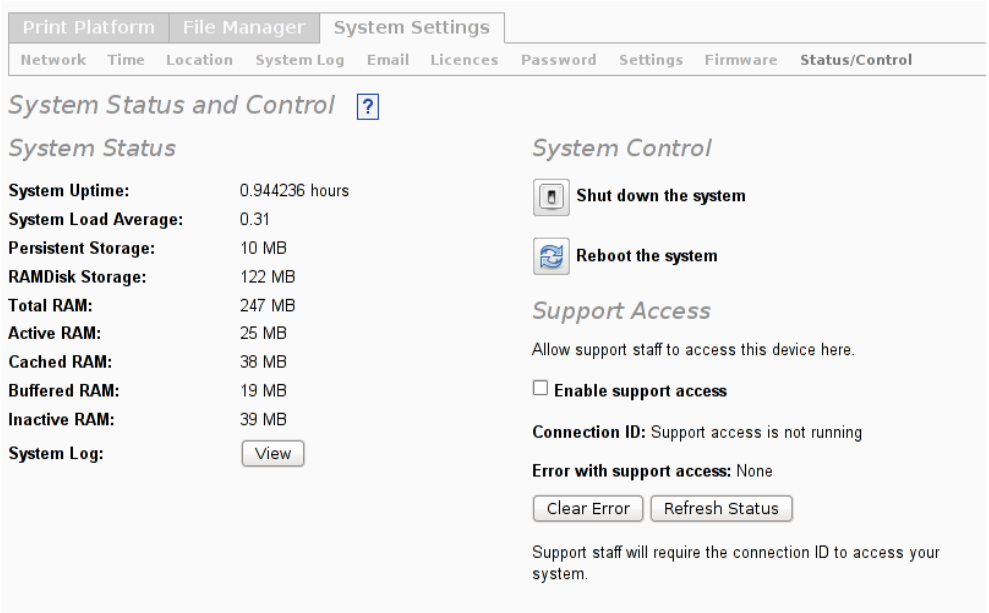

System Status

This section shows the current system status. Most of the entries in the list are self-explanatory.

"Persistent Storage" is the amount of storage available on the on-board hard disk or flash disk module for storing user files such as PDF forms. When this runs low, the user will be unable to add new PDF forms to the appliance.

The "RAMDisk Storage" is used for all temporary files and is half the size of the total installed RAM. If this runs low, then print jobs may stop processing through the PdfPlus. During the implementation phase, it is important to ensure that the maximum print load will not cause the RAMDisk Storage to run out.

The "View" button for the "System Log" will display a pop-up window showing the last few lines of the internal system log (which resides on the RAMDisk and is not persistent across reboots).

The "View" button for "Network Mounts" displays a window showing mounted network shares. This is used to verify that any archive and hot folder shares which have been con-

figured have been successfully mounted by the appliance and have sufficient storage available.

System Control

This page provides shut down and reboot buttons.

Shut down

The shut down button will display a confirmation dialog box. After confirmation has been given, the system will go through a safe shut down procedure. The web interface will display a success message when it is no longer able to connect to the appliance's embedded web server.

Reboot

If the user confirms a reboot, the appliance will safely shut down, then reboot. The web interface will re-connect to the appliance login screen as soon as the reboot is complete.

Support Access

The PdfPlus provides a remote access mechanism to allow support staff to access the appliance and investigate problems.

To use this feature, the appliance must have internet access. If a firewall is in use, it must be configured to allow the appliance to initiate a direct connection on TCP port 80 (HTTP) to the support access server (rssh.wmltd.co.uk). This is equivalent to allowing the appliance to browse the web.

To turn support access on, check the box “Enable support access”. The appliance will attempt to connect to the support access server. When a successful connection is made, a connection ID will be displayed. Support staff will request this connection ID to enable them to connect back to the appliance.

Chapter 5. Queue Management

Queue management is carried out in the **Print Platform** section of the PdfPlus user interface. This section contains the queue management interface, comprising the **Queues** page, which lists and monitor all queues on the appliance, and the **Queue Config** section, where detailed configuration of queues can be carried out.

Other pages in this section are the **Jobs** page, on which recent jobs are listed and managed, and the **Platform Settings** page, on which platform-wide settings may be configured.

5.1. Queue management interface

Simple queue administration tasks, such as adding a queue or changing input and output details, may be carried out on the **Print Platform > Queues** page. Full configuration is carried out in **Print Platform > Queue Config**, described in Section 5.2.

The first point of entry to the queue management interface is the print queues list (Figure 5.1), found on the **Print Platform > Queues** page.

Figure 5.1. The print queues list

This list displays and monitors all queues on the PdfPlus. The summary view shown in Figure 5.1 displays a status line for each queue, comprising the queue name, activity monitor and queue status information, as well as a number of print user actions for the queue. The queue list is updated periodically so that the information shown always reflects the current state of the queues.

System status information is displayed above the queue list. The system load average and the total and inactive memory are displayed. The information is updated periodically to show the current values.

Figure 5.2 details the components of the queue status line.

Figure 5.2. The queue status line

❶ Queue info button. Click to display a summary of the queue configuration.
❷ Activity monitor. The bars light up when jobs are processing on the queue.
❸ Queue status. The current status of the queue and information about the last job processed on the queue.
❹ Print user actions: print file, view jobs and print status page.

Print user actions

The queue status line enables a number of print user actions, which are described below.

Print file

Click on the print icon ( ) to send a file to print on the queue. You will be prompted to browse to a file. Once a file has been selected, it will be sent to print.

View jobs

The view jobs icon ( ) provides a direct link to the **Print Platform > Jobs** page (see Section 5.4), set to display all jobs on the selected queue.

Print status page

Click on the print status icon ( ) to send to print a queue statue report page (Figure 5.3) providing queue details and status information and current system information.

Figure 5.3. Queue status report page

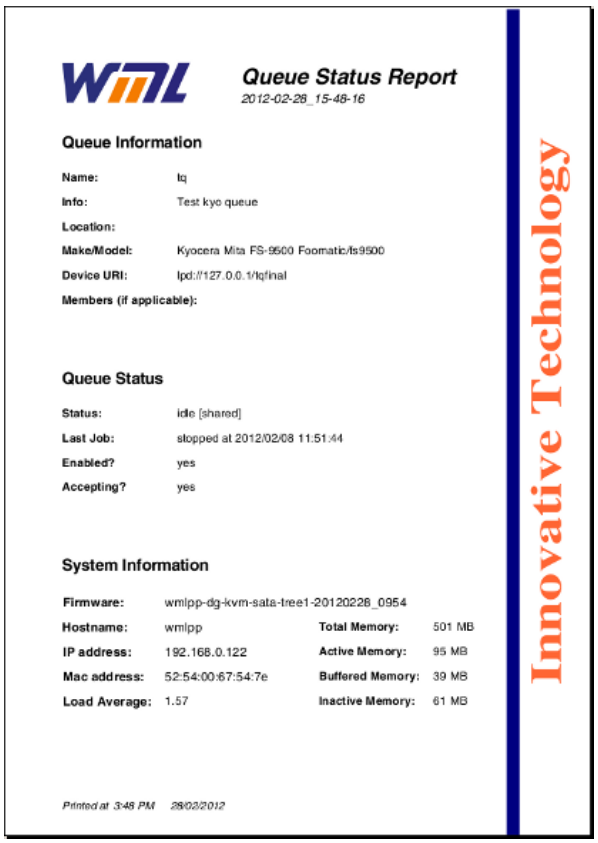

Queue Status Report
2012-02-28_15-48-16

Queue Information

Name: lq
Info: Test kyo queue
Location:
Make/Model: Kyocera Mita FS-9500 Foomatic/fs9500
Device URI: lpd://127.0.0.1/lqfinal
Members (if applicable):

Queue Status

Status: idle [shared]
Last Job: stopped at 2012/02/08 11:51:44
Enabled? yes
Accepting? yes

System Information

Firmware: wmlpp-dg-kvm-sata-tree1-20120228_0954
Hostname: wmlpp
IP address: 192.168.0.122
Mac address: 52:54:00:67:54:7e
Load Average: 1.57
Total Memory: 501 MB
Active Memory: 95 MB
Buffered Memory: 39 MB
Inactive Memory: 61 MB

Printed at 3:48 PM 28/02/2012

Innovative Technology

Print admin controls

The print queues list provides admin controls which allow simple administrative tasks to be carried out on existing queues and which facilitate the addition of new queues.

The print admin controls for an existing queue are revealed by clicking on the status line for the queue. The controls will appear in a new line below the queue status line, and labels for the controls will appear in the queue list header (Figure 5.4). Click on the status line again to hide the controls.

> **Tip**
>
> To show or hide the print admin controls for all queues, click on the queue list header.

Figure 5.4. Print admin controls

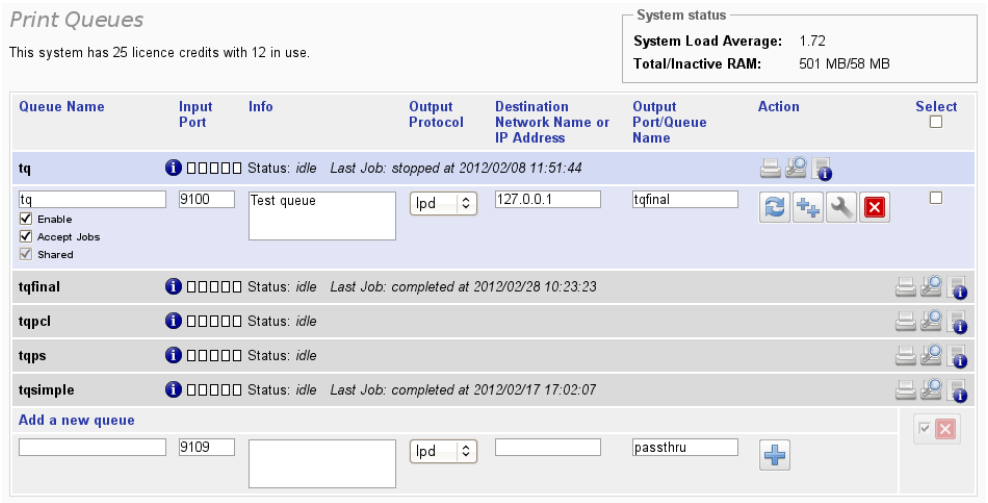

> **Note**
>
> Please note that **Info** field is hidden at narrow browser widths. For access to this field, your browser window should be 1000px or more wide.

Editing a queue

Changes made to the queue name take effect as soon as focus leaves the queue name field. Modifications made to other fields should be saved by clicking the update button ( ).

> **Note**
>
> The queue name may contain upper and lower case letters, numbers, the dash character and the underscore character. Queue names may not contain spaces.

Click on the configure button ( ) to gain access to the full queue configuration interface for the queue.

A queue may be duplicated by clicking on the duplicate button ( ). This will create a copy of the queue, including full configuration details. You will be given the opportunity to edit the automatically-generated name for the new queue before the new queue is created.

An existing queue may be removed by clicking on the delete button ( ). You will be prompted to confirm the deletion. You can also delete multiple queues by checking the "Select" checkbox for each queue, then clicking on the delete selected button ( ).

Adding a print queue

To add a queue, fill in the queue name, input port, output protocol, destination network name or address and output port/queue name in the "Add a new queue" area, then click on the add button (). The queue will be added to the system.

> **Note**
>
> The queue name may contain upper and lower case letters, numbers, the dash character and the underscore character. Queue names may not contain spaces.

The input port defaults to the next available port number, starting at 9100, a value which is commonly used in printing applications. It is possible for two queues to respond on the same raw port. In this case the first queue (alphabetically) will respond to incoming print data on the shared port, **unless you have enabled per-IP address filtering in the input tab of the queue configuration**. To find out more about how queues can respond based on sender and destination IP addresses, see the section called "LPD/raw input settings".

Classes

The PdfPlus can be configured with *print classes*. A print class is a group of print queues. Print data sent to a print class will be forwarded to members of the class in rotation.

The classes interface is similar to that for print queues (Figure 5.5).

Figure 5.5. The classes list

To add a class, enter a name for the class, and select at least one member for the class. Multiple queues may be selected by holding down the control key. The input port defaults to the next available port number, starting at 10100. Click on the add button () to add the class to the system.

> **Note**
>
> At least one queue must be configured before it is possible to add a class.

Click on the configure button ( ) to gain access to the full configuration interface for the class. Classes have **Input** and **Output** options, but no data processing options; all data processing is configured on a per-queue basis.

5.2. Queue configuration

While some queue parameters can be modified in the print queue list, full configuration options for each queue are reached via the configure button in the print queues list, as shown in Figure 5.6, or by going to the **Print Platform > Queue Config** page and switching to the required queue using the **Switch to queue** menu.

Figure 5.6. The configure button in the print queues list

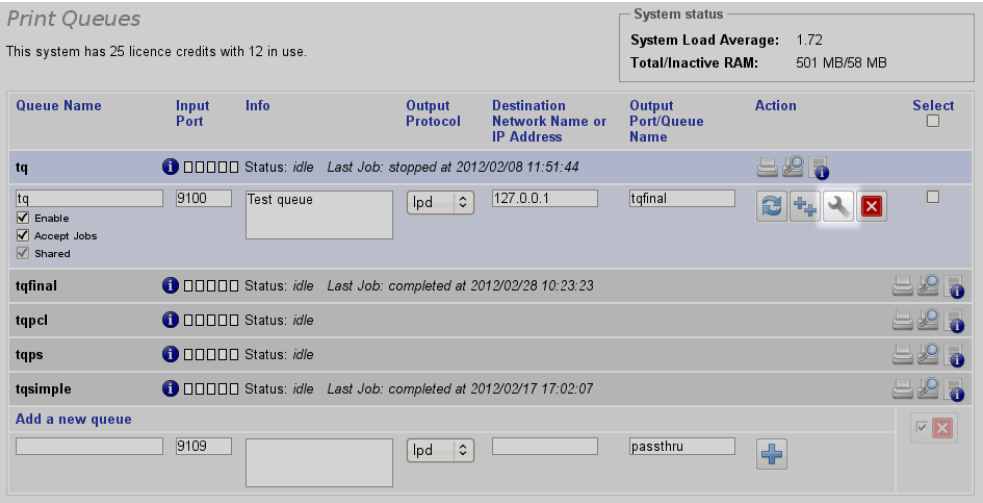

The following sections describe the options available for configuration within the **Print Platform > Queue Config** section.

Data input

Print data can be sent to the PdfPlus using a variety of methods. These include the well-known Berkeley *LPD protocol*, defined in Internet RFC 1179, the *IPP protocol*, which is defined in RFC 3510 and the *raw print protocol*, by which data is sent unmodified, with no additional command metadata, to a particular TCP port.

In addition to these input protocols, the PdfPlus can can receive data from an *HTTP POST request*, where the print data is encoded in a POST form and it can also collect print files from a network-shared hot folder. This is called the *input share* or *hot folder* method.

The PdfPlus also has a unique feature which allows it to behave as if it were any number of printers on a network, each using the same TCP port for raw printing. The PdfPlus can be configured with a number of IP address aliases. Any queue can then be instructed to "listen" only on a particular IP address, ignoring input to other IP addresses.

Select the **Input** tab (Figure 5.7) to configure the input data options for the queue.

Figure 5.7. Queue configuration: Input tab

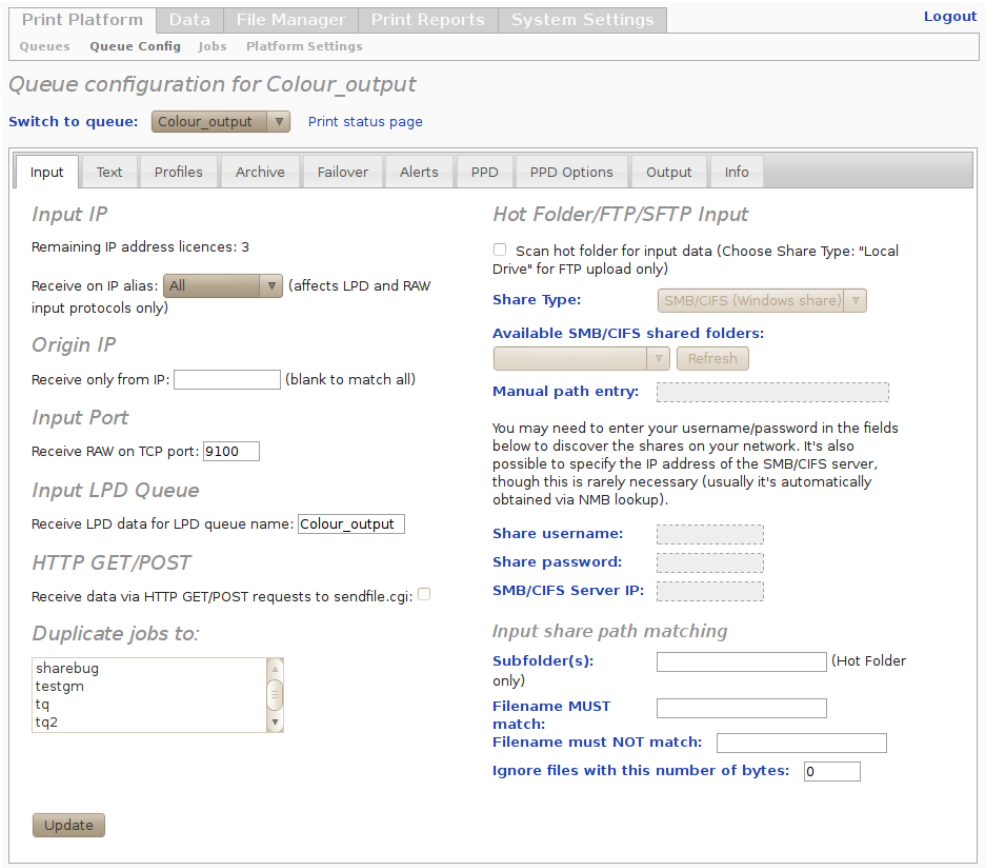

The left-hand column contains the input settings for LPD/raw input data. The right-hand column contains the input share settings.

> **Caution**
>
> Before changing any settings, shut down any print job scheduler which may be sending data to the PdfPlus. If print jobs are spooling onto the PdfPlus when the queue settings are updated, data loss may occur.

LPD/raw input settings

The input IP on which to listen for raw print data can be selected from the “Receive on IP alias” drop-down control. This shows the available IP addresses on which the device can

listen, including all the aliases. To add aliases, use the Network settings page in **System Settings > Network**.

It is possible to filter jobs using the *origin IP*. Enter an IP address in the **Receive only from IP** field and the queue will only accept jobs originating from that single IP address.

> **Note**
>
> It is only possible to specify a single IP address or "all" for the input IP and the origin IP. To process data from two or more originating servers, create additional print queues with identical settings and set input/origin IP settings for the two or more addresses.

The **Input Port** for raw TCP data input can be set to a value in the range 1024 to 32767, inclusive.

> **Note**
>
> The input port can also be set in the Print Queues list.

> **Note**
>
> If the CDL module has been licenced (DataGateway only) some additional fields will be visible here, allowing you to specify that the incoming data is expected to be CDL-formatted data.

The **Input LPD Queue** field can be set to a value with less than 256 characters and no spaces. This alias can be used with the LPD input protocol in place of the actual name of the queue. The LPD queue alias defaults to the name of the queue.

HTTP GET/POST input

To enable the queue to receive data via an HTTP GET or POST request to **sendfile.cgi**, check the box in **HTTP GET/POST**.

> **Note**
>
> The HTTP GET/POST input feature uses a licence credit for each queue on which it is enabled.

Making an HTTP POST request

> **Note**
>
> While it is possible to use an HTTP GET request to send data to the PdfPlus, it is more natural to use HTTP POST, which is the recommended request to use.

To print some data to the PdfPlus, make an HTTP POST request with the following parameters in POST form (use the application/x-www-form-urlencoded POST request type):

- queue: The queue to which you wish to send the data (required).
- copies: The number of copies of the data which should be output (defaults to 1 if omitted).
- title: The title for the job (optional).
- user: The user who sent the data (optional).
- data: The data for the job (required).

All parameters should be url-encoded. This will be necessary for most of the fields; certainly for the data.

Make the request to **http://your.PdfPlus.net/sendfile.cgi**. An XML-formatted response will be returned.

A useful utility for testing is *curl*, available, at the time of writing, from http://curl.haxx.se/.

The following call to the curl command will spool one copy of the file "file.name" giving it the title "test", and printing it to the queue "testgm" on the PdfPlus "wml.localnet" on behalf of the user "Bob":

```
curl -X POST http://wml.localnet/sendfile.cgi -d queue=testgm -d copies=1 \
-d title=test -d user=Bob --data-urlencode data@file.name
```

> **Note**
>
> If copies is set to a number N, greater than 1, then the data will be spooled *as input* N times, rather than having the output of any processing on the queue emitted N times. Any processing which the PdfPlus would carry out on the data will be repeated N times.

The response to an HTTP POST request

The response returned by call to **sendfile.cgi** is XML data which indicates the success or failure of the request.

If the request was successful, the XML will contain an element called **success** meaning that the data was accepted onto the queue. It doesn't mean that the data was successfully *output*, only that the data was received and accepted onto the queue.

If there was an error which prevented the data from being accepted, then a field called **error** will be found in the XML response. This error message will describe what went wrong.

In addition to the success and error elements in the XML, the parameters of the request will be returned back to the caller, which can be helpful when testing the process.

An example of a successful response is shown below. The **success** element contains the message "Successfully spooled data to" followed by the target queue name (testgm, in the example).

```
<?xml version="1.0" ?>
```

```
<!DOCTYPE WmlSendFileResponse>
<WmlSendFileResponse>
  <copies>1</copies>
  <dataSizeBytes>674</dataSizeBytes>
  <jobId>34933</jobId>
  <queue>testgm</queue>
  <success>Successfully spooled data to testgm</success>
  <title>test</title>
  <user>seb</user>
</WmlSendFileResponse>
```

In the next response, the call failed because the user did not pass a data file:

```
<?xml version="1.0" ?>
<!DOCTYPE WmlSendFileResponse>
<WmlSendFileResponse>
  <dataSizeBytes>0</dataSizeBytes>
  <error>Please pass document data in the parameter "data".</error>
</WmlSendFileResponse>
```

The full list of elements in the response XML

- copies: The number of times to spool the data to the target queue.
- dataSizeBytes: The number of bytes of data presented in the POST (or GET). Always returned, whether request is a success or not.
- error: An error message. If present, the data was not accepted onto the queue.
- jobID: The internal CUPS job ID for the successfully spooled data. Only present if the job is successfully spooled. If copies was set to >1, then this will contain a comma-separated list all IDs.
- queue: The target queue specified in the request.
- success: A success message. Only present if data is successfully accepted.
- title: The job title specified in the request.
- user: The user specified in the request.

The most common HTTP POST request error messages

- Please pass document data in the parameter "data". *Check your HTTP request.*
- The target queue does not accept HTTP GET/POST requests to spool data. *Make sure you have enough licence credits and ensure that HTTP GET/POST is enabled in **Print Platform > Queue Config > Input** .*
- Please set the target queue in the parameter "queue". *Check your HTTP request.*
- The queue is not accepting jobs. *Check that the **Accept** box is checked for your target queue in **Print Platform > Queues**.*
- The print transport (CUPS) did not accept the job. *Check you haven't used up all the ramdisk or persistent storage on the system (**System Settings > Status/Control**). Try re-starting CUPS at **Print Platform > Platform Settings > Cups Control**.*

Duplicating input jobs

The **Duplicate jobs to:** box is a multi-select menu. Any queues highlighted in this box will receive a copy of every input data file received on the current queue. Hold the Ctrl key to select multiple queues in this box.

Share-based input

To print files from a network share, the PdfPlus first "mounts" the network share, then scans it on a regular basis (every second) for new files to print. When it finds a file, it monitors its size, then when the size has been static and non-zero for 3 seconds (the "settle-time"), it spools it for print.

Empty files (with zero size) are not printed. The PdfPlus logs the existence of empty files in the input share with a message like this:

```
Oct 16 10:08:35 wmldevice ctrl[731]: \
           Ignoring zero sized file: '/tmp/hotfolder/myqueue/file.pdf'
```

This log message is repeated every 60 seconds.

> **Caution**
>
> The input share settle-time is only *3 seconds*.
>
> Files *must* be fully copied onto the input share within this time-scale.

Care must be taken with the share-based input method if your application creates files on the input share using file access programming techniques.

If you use functions such as fopen() in C or file-streams in C++, you must ensure that your application creates the file in less than the settle-time of 3 seconds.

Suppose your application creates a file in the input share, sends *some*, but not *all* of the print data into the file, then pauses for more than the settle-time before completing. The PdfPlus will see the the file size is non-zero and has remained static for 3 seconds, so it will pick up the incomplete job, and attempt to print it.

To avoid this problem, you can use two strategies: Either have your application generate the file in a separate folder and then copy the completed file into the PdfPlus input share *or* have your application create each file with a ".spl" suffix (meaning SPooLing). The PdfPlus will ignore files with a ".spl" suffix (which must be in lower case). When your application has completed generating the file, have it re-named without the .spl suffix. The job will then be printed.

Once printed, the PdfPlus erases the file from the network share. For this reason, the network share used for printing should be configured for read/write access and if a copy of the print data is required it should either be archived using the PdfPlus's archiving features, or archived before it is dropped in the input share.

To enable share-based input, check the box labelled **Scan share directory for inputs**. It's possible to use either Unix/Linux Network File Server shares or Windows/Samba SMB/CIFS network shares from which to collect print data.

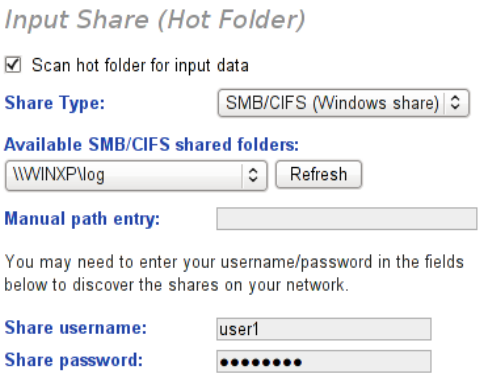

Selecting a share

NFS (Linux/Unix shares)

To connect to a Linux/Unix NFS share, choose **NFS (Linux/Unix)** from the **Share Type** drop-down. Enter the path to the NFS share in the following format:

`[Network/IP Address]:[folder path]`

For example:

`192.168.100.100:/input_data`

The NFS server's exports file should be configured to allow both read and write access to the directory path (/input_data in the example). An example of a suitable line in /etc/exports on the NFS server is:

```
/input_data 192.168.0.0/255.255.255.0(rw,no_root_squash,sync,subtree_check)
```

If the NFS share is to be used as an *archive* Share, then correct permissions must be set on the share directory on the NFS server. The following would be required for an archive share specified as:

`192.168.100.100:/archiveData`

On the server it would be necessary to change the user on this directory to "root" or 0 and the group of the directory to 7, which is the GID for the "lp" user on some systems:

`root@nfsServer# chown 0:7 /archiveData`

Finally, ensure that both user and group are able to read and write to the directory:

`root@nfsServer# chmod ug+rwx /archiveData`

SMB/CIFS (Windows shares)

For Windows SMB/CIFS network shares, select **SMB/CIFS (Windows share)** and click on the **Refresh** button to scan the network for available shares. This will populate the **Available SMB/CIFS shared folders** drop-down.

> **Important**
>
> *Most recent Microsoft Windows file servers will not broadcast their available shares to the network without authentication.*
>
> If necessary, enter your network username and password in the fields **Network user name** and **Network password** and then press **Refresh**. The username and password will be used to obtain the list of available shares.

It is also possible to enter the share path manually in the **Manual path entry** text field, by choosing the option **Enter path in text box** from the **Available SMB/CIFS shared folders** drop-down.

Enter the SMB/CIFS path in the format:

`\\MACHINE\shareName`

The **Network user name** and the **Network password** fields must be filled in to attach shares which require authentication.

Unsupported share types

The PdfPlus print platform uses a subset of the tools from Samba (www.samba.org) to scan for, and attach, SMB/CIFS network shares. Some Microsoft Windows share configurations are not supported and cannot be used for input share access on the PdfPlus. In particular, Active Directory shares cannot be used. There may also be problems with certain WINS configurations.

Overlay templates and logic profiles

The profiles tab allows you to “bind” an overlay template or a logic profile to the queue.

An overlay template is a series of instructions to be applied to the print data passing through a queue, such that the resultant print job may be overlaid with a set of fixed or variable data texts, or the contents of one or more PDF files.

A logic profile is a set of rules to apply to the data stream. Logic profiles scan the text in the data stream and apply “logic actions” where conditions are matched.

To change the overlay template for the queue, select the **Profiles** tab (Figure 5.8).

Figure 5.8. Queue configuration: Profiles tab

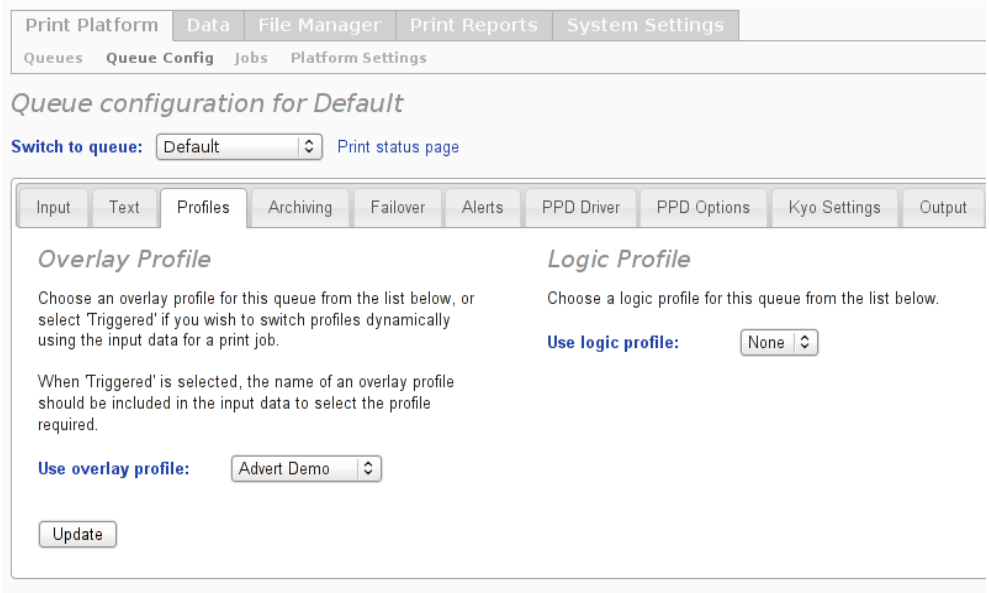

You may choose to apply an existing overlay template to the queue by selecting the template name from the drop-down control. Alternatively, if you wish to switch overlay set dynamically using the input data for a print job, you can select **Triggered**. The print input data should then include the name of the overlay template to be used.

Logic profiles are bound to the queue in a similar way using the corresponding drop down menu of available logic profiles.

> **Note**
>
> If both an overlay template *and* are bound to a queue, then any overlay template which is triggered by the logic profile will be applied to the job. If no template is triggered by the logic profile, the template selected here will be applied.

Configuration of overlay templates and details about the structure of input data are described in Chapter 6. Logic profiles are described in Chapter 7.

Archiving

A queue may be configured to archive print jobs to network attached storage. The network share may be either a Unix/Linux NFS share or a Windows/Samba SMB/CIFS network share.

To configure the archive settings for the queue, select the **Archive** tab (Figure 5.9).

Figure 5.9. Queue configuration: Archive tab

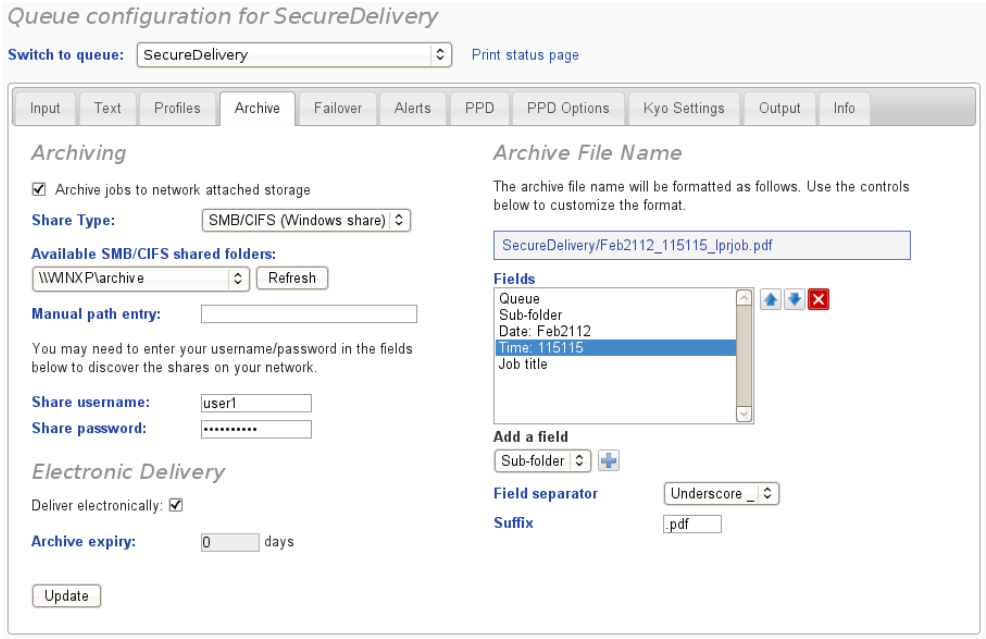

To turn on archiving, check the **Archive jobs to network attached storage** box. This will activate the **Share Type** drop-down and other form controls, allowing you to select the storage to be used for the archive.

Selecting a share

NFS (Linux/Unix shares)

To connect to a Linux/Unix NFS share, choose **NFS (Linux/Unix)** from the **Share Type** drop-down. Enter the path to the NFS share in the following format:

`[Network/IP Address]:[folder path]`

For example:

`192.168.100.100:/input_data`

The NFS server's exports file should be configured to allow both read and write access to the directory path (/input_data in the example). An example of a suitable line in /etc/exports on the NFS server is:

```
/input_data 192.168.0.0/255.255.255.0(rw,no_root_squash,sync,subtree_check)
```

If the NFS share is to be used as an *archive* Share, then correct permissions must be set on the share directory on the NFS server. The following would be required for an archive share specified as:

`192.168.100.100:/archiveData`

On the server it would be necessary to change the user on this directory to “root” or 0 and the group of the directory to 7, which is the GID for the “lp” user on some systems:

```
root@nfsServer# chown 0:7 /archiveData
```

Finally, ensure that both user and group are able to read and write to the directory:

```
root@nfsServer# chmod ug+rwx /archiveData
```

SMB/CIFS (Windows shares)

For Windows SMB/CIFS network shares, select **SMB/CIFS (Windows share)** and click on the **Refresh** button to scan the network for available shares. This will populate the **Available SMB/CIFS shared folders** drop-down.

> **Important**
>
> *Most recent Microsoft Windows file servers will not broadcast their available shares to the network without authentication.*
>
> If necessary, enter your network username and password in the fields **Network user name** and **Network password** and then press **Refresh**. The username and password will be used to obtain the list of available shares.

It is also possible to enter the share path manually in the **Manual path entry** text field, by choosing the option **Enter path in text box** from the **Available SMB/CIFS shared folders** drop-down.

Enter the SMB/CIFS path in the format:

```
\\MACHINE\shareName
```

The **Network user name** and the **Network password** fields must be filled in to attach shares which require authentication.

Archive file name format

The format for the file names of archived print jobs can be controlled in the **Archive File Name** section of the **Archive** tab. An example file name is shown in the box, and the controls below can be used to customize the file name format.

The **Fields** box displays the list of fields used to make up the current archive file name. Add fields by choosing them from the dropdown menu and clicking on the add button ( ). You can re-order or delete fields by highlighting them and clicking on the up and down arrows or delete button. You can also choose the field separator character and the file suffix to be used. The example file name will be updated to reflect your changes.

To save your changes, click on the **Update** button.

Archive expiry

You can choose to have archive files "expire" after a given number of days. This means that they will be removed from the system after the given time period has elapsed. If the **Archive expiry** value is set to 0, then archive files will not expire.

Failover

The PdfPlus can be placed into *failover mode* in **Platform Settings > Failover**. In failover mode, incoming print jobs will be passed to the "failover queue(s)".

To set the failover queues for the current queue, select the **Failover** tab (Figure 5.10).

Figure 5.10. Queue configuration: Failover tab

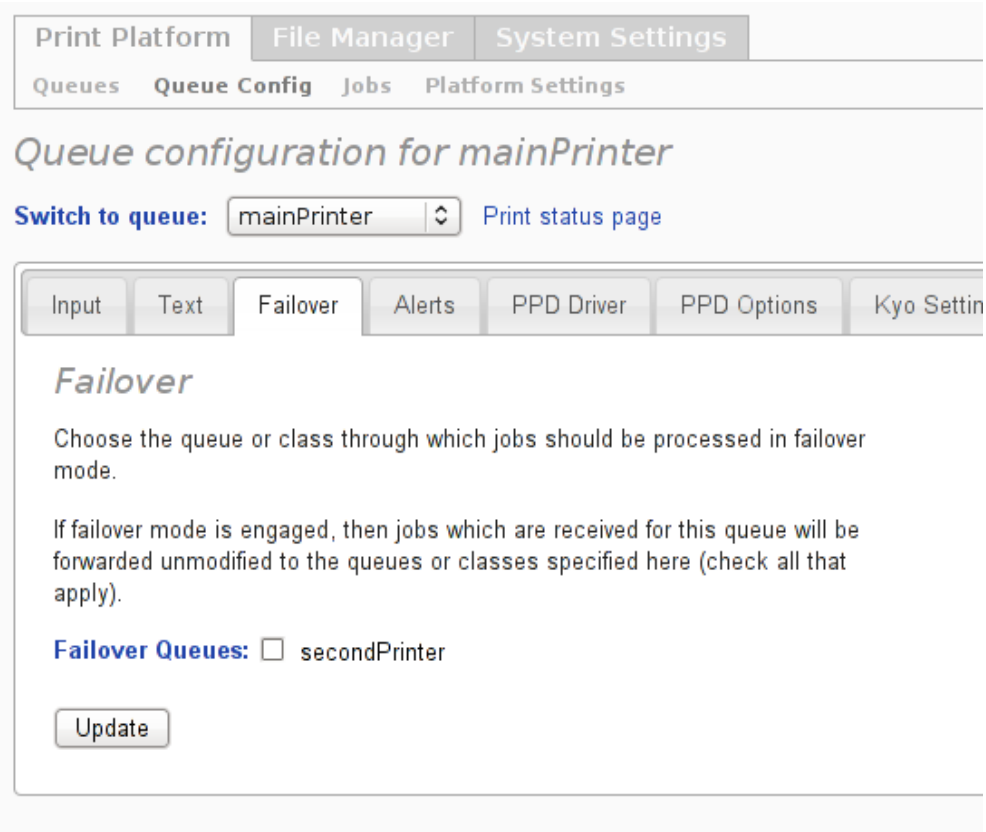

You may select the queues through which jobs should be processed in failover mode.

If failover mode is engaged, then jobs which are received for the current queue will be forwarded unmodified to the queues specified.

If **Pool** is selected, then jobs will be failed over to each of the selected queues in turn.

To save your changes, click on the **Update** button.

Alerts (Queue Monitoring)

The PdfPlus system may be configured to send alerts if the queue does not appear to be processing print jobs at the expected rate. The alerts are sent as emails, using the main system email settings configured in **System Settings > Email**, or per-queue settings, if provided.

Alerts can be sent when too few documents are processed, on a regular basis with a status message about how many jobs have been processed, or for every individual document processed on the queue.

To configure the queue alert settings, select the **Alerts** tab (Figure 5.11).

Figure 5.11. Queue configuration: Alerts tab

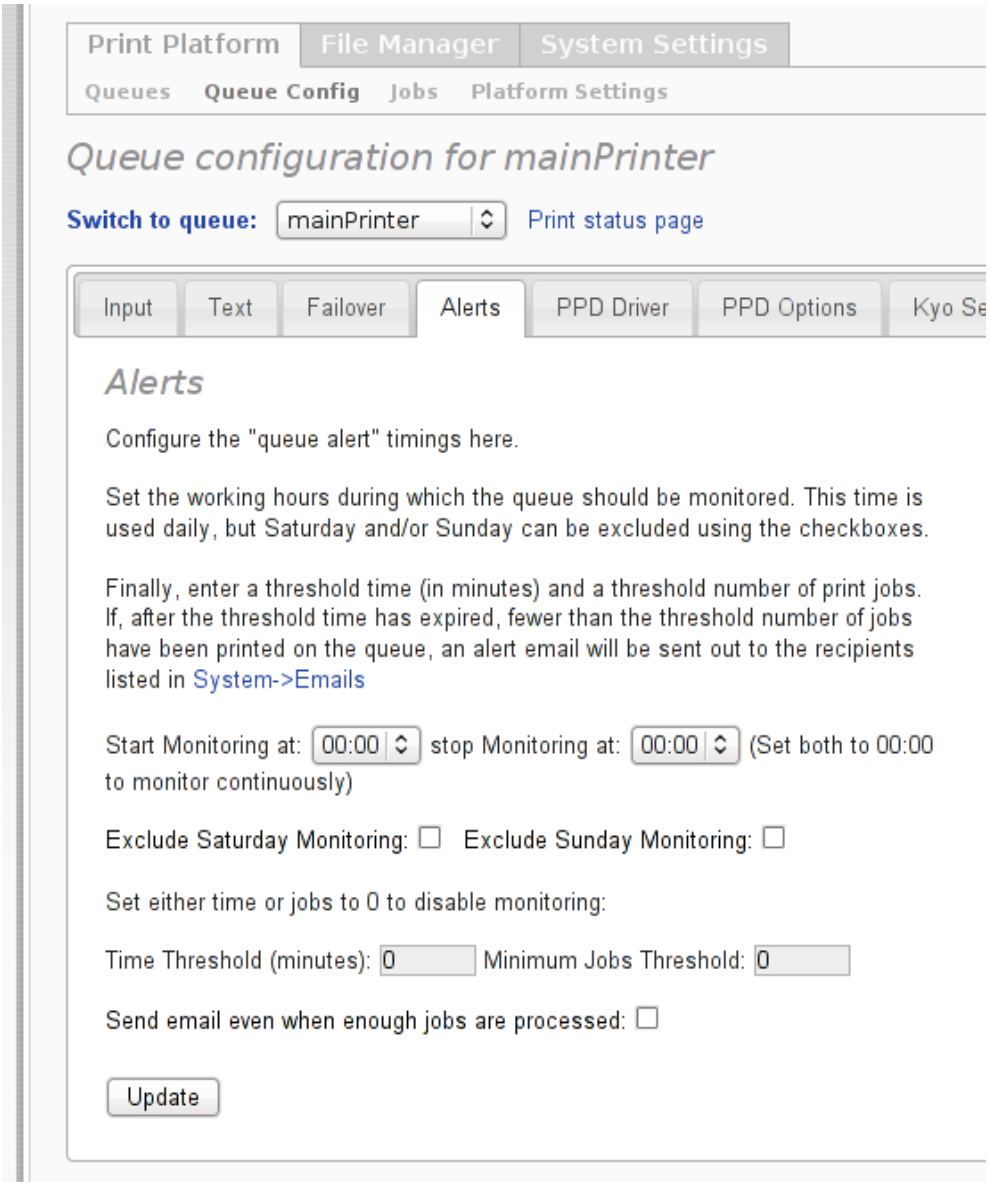

Firstly, set the working hours during which the queue should be monitored using the drop-down controls. Choose a start time and an stop time for the queue monitoring period. This

monitoring period will be applied daily, though Saturday and Sunday can optionally be excluded using the checkboxes for five and six day production weeks.

To monitor the queue continuously, set both **start monitoring** and **stop monitoring** times to **`00:00`**.

Next, enter a threshold time (in minutes) and a threshold number of print jobs. If, after the threshold time has expired, fewer than the threshold number of jobs have been printed on the queue, an alert email will be sent out to the recipients listed in **System Settings > Email** (see "Email settings" in Chapter 4: *PdfPlus Configuration*).

Set either **Time Threshold** or **Minimum Jobs Threshold** to **`0`** to disable monitoring.

If a status message is required even when the minimum threshold has been exceeded, check the **Send monitoring email even when...** checkbox.

To configure the system to send an alert email for every single job processed during the monitoring working hours, check the box **Send monitoring email for EVERY job...**.

To save the queue alert settings, click on the **Update** button.

PPD driver

To select a PPD driver for the queue, select the **PPD Driver** tab (Figure 5.12).

Figure 5.12. Queue configuration: PPD Driver tab

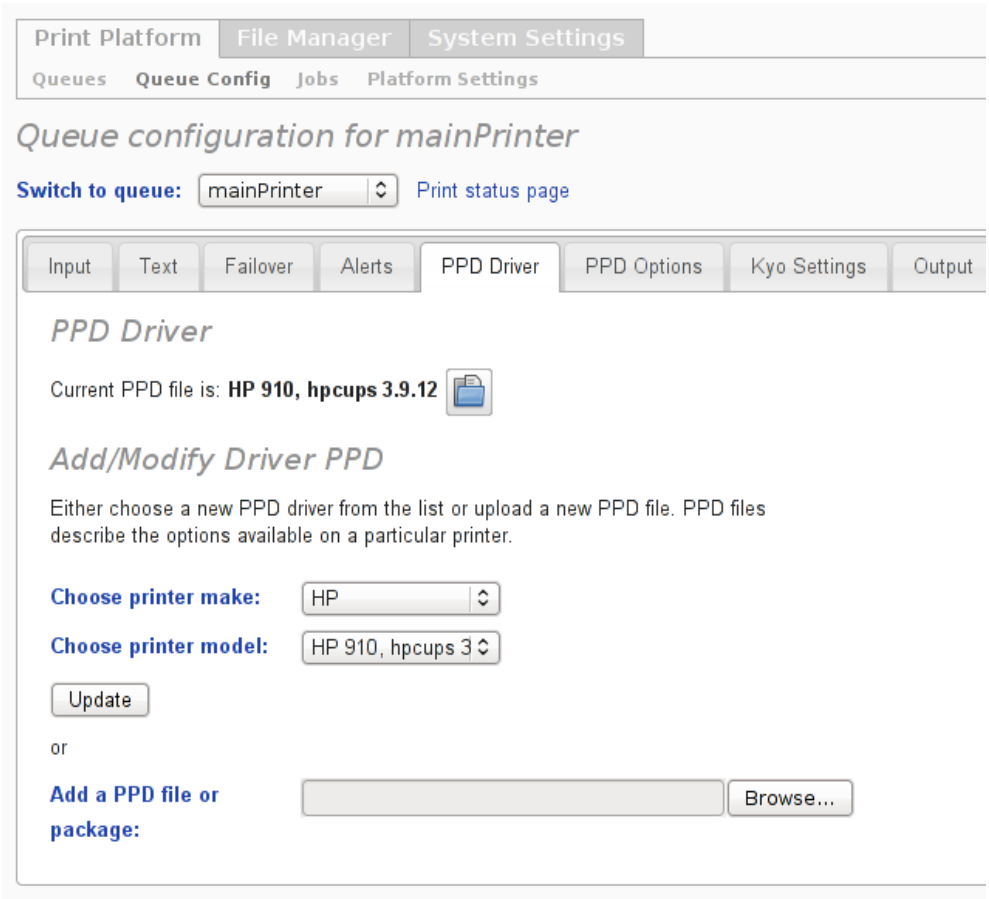

A PPD file is a *printer description*. It describes the options available for a particular printer and what data format the printer requires (PostScript, PCL, raster data etc). It's not actually a printer driver, but it does specify which driver to use, and what the parameters of the print driver should be.

The current PPD file is displayed at the top of the page. You can view the content of the file by clicking on the file name.

A selection of PPD files are shipped on the PdfPlus. These are available for selection in the “Choose printer make” and “Choose printer model” menus. First choose the make of your printer, then a selection of models will be listed in the “Choose printer model” menu.

It is also possible to upload an individual PPD file or a package of PPD files via the “Add a PPD File or package” file upload dialog.

If an individual PPD file is uploaded, it will be stored (and added to the make/model menus) and set as the current PPD.

A package of PPD files is added to the make/model menus. PPD packages can be rpm packages, deb packages, zip archives or tar archives. These are automatically uncompressed and unpacked.

All uploaded PPD files are listed in **File Manager > PPD Files**

When a PPD file is set for a queue, the driver options for the printer will be made available (see “Driver options”).

Internal data PPD files

The PdfPlus contains PPD files which force print data output to be PDF, generic PCL or “raw”, where not modification is made to input print data. These PPDs all have the make “WilliamMatthew”.

Driver options

To modify printer driver options for the queue, select the **PPD Options** tab.

Printer driver options are displayed here. The options will match the current PPD file (see “PPD driver”).

Figure 5.13. Queue configuration: Driver Options tab

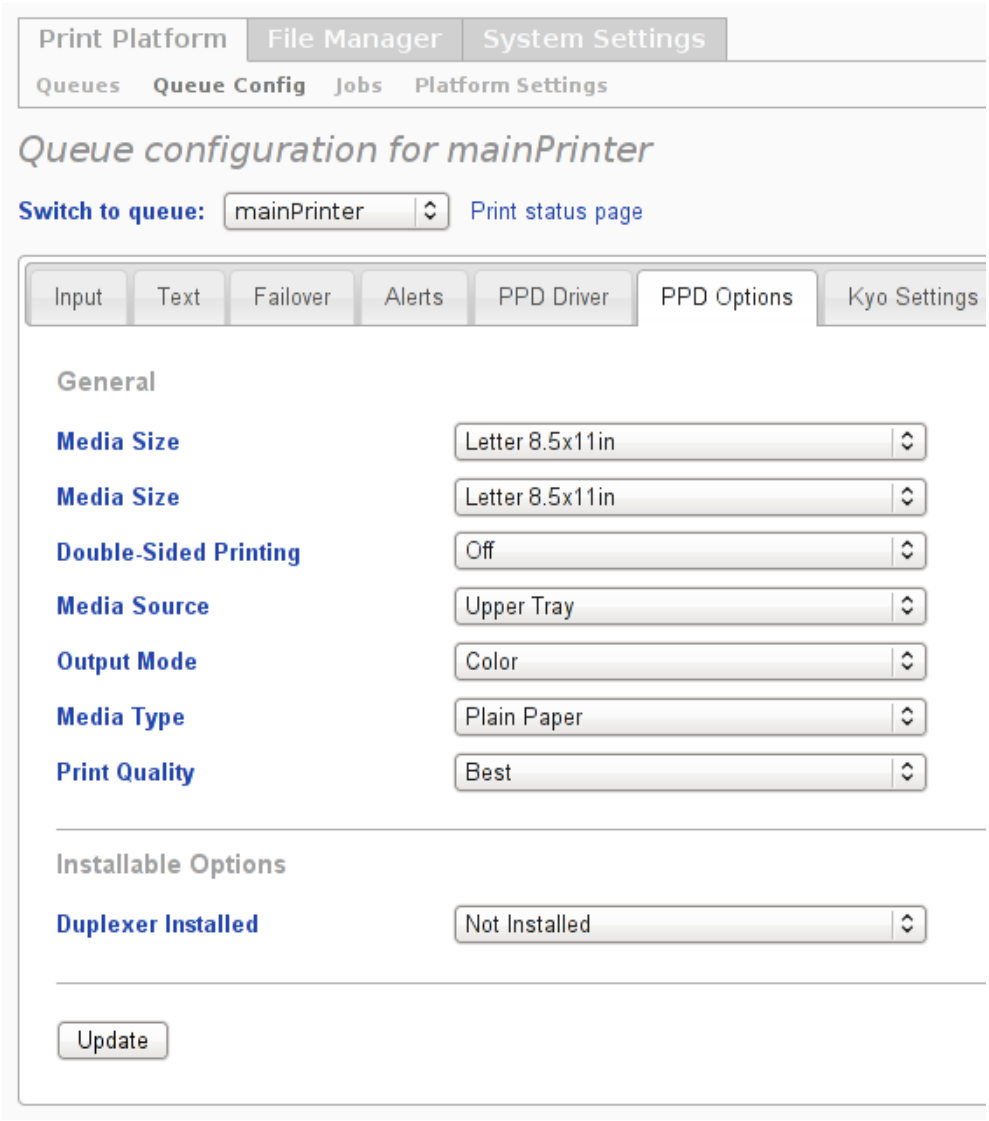

Use the drop-down controls to modify the driver options as required. To save your changes, click on the "Update" button at the bottom of the list of options.

Output

The output options for a queue are configured in the **Output** tab (Figure 5.14).

Figure 5.14. Queue configuration: Output tab

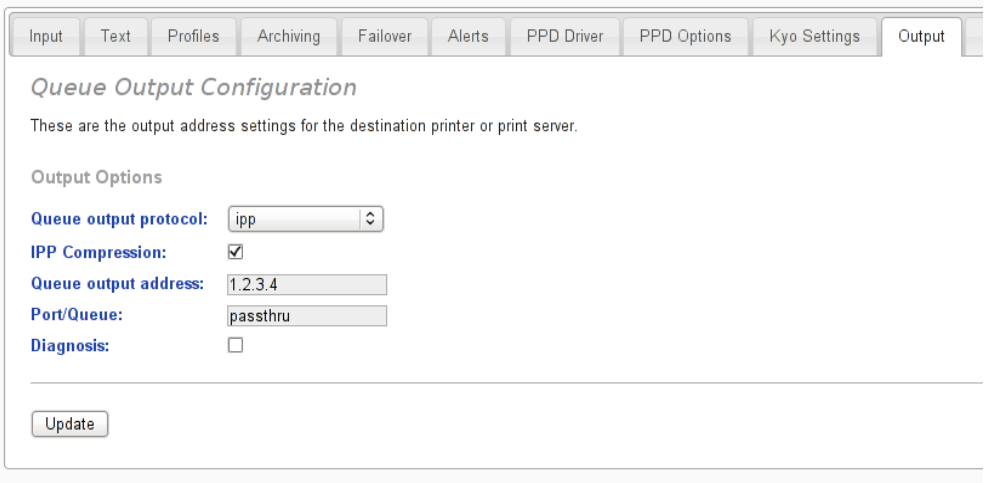

Common output options

For lpd, raw, ipp, and ipps protocols, select the required output protocol from the menu and enter a suitable destination address (Queue output address) and port or queue.

IPP options

If either “ipp” (Internet Printing Protocol) or “ipps” (Secure IPP) is chosen for “Queue output protocol”, it is possible to specify whether IPP compression should be applied to outgoing data. Use of this feature requires that the receiving IPP device is able to decompress the data. If the receiving IPP device is not able to apply decompression, then the data will be sent without compression.

Generally IPP print jobs are sent to port 631. If a different port is required, do not enter it in “Port/Queue” - this is where the IPP destination queue name must be entered. Instead, enter the port with the address in the following format:

```
192.168.0.100:632
```

Queue diagnostics

For diagnostic purposes, the print job can be saved before and after it is processed by the PdfPlus software. The input data will appear in the file list as prefilter_queue.bin and the output as postfilter_queue.bin.

Enable diagnostics by checking the “Diagnosis” checkbox.

This feature has the potential to fill the storage of the device and lead to failed print jobs, so care must be taken with its use. Because of this, the feature will always be switched off following a reboot of the PdfPlus.

Queue info

The queue info tab (Figure 5.15) displays a summary of queue and printer details for the current queue.

Figure 5.15. Queue configuration: Queue Info tab

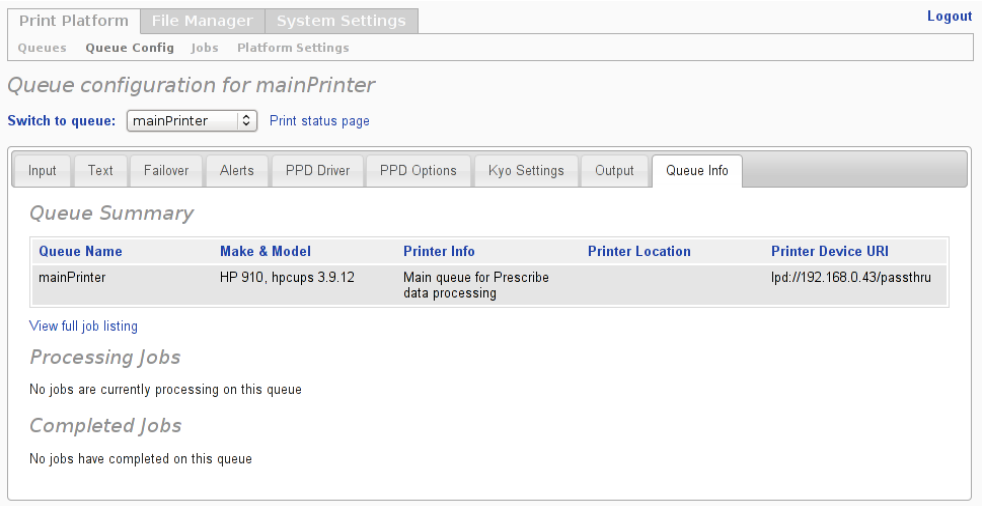

5.3. Queue configuration: Example workflow

This section describes the basic steps required to create a new, usable print queue on the PdfPlus. It doesn't describe every setting (refer to Section 5.2 for that), but instead gives an overview of the initial configuration of a new queue.

First go to **Print Platform > Queues**. To create a new queue, we will fill in the details in the "Add a new queue" row of the "Print Queues" table.

Figure 5.16. Setting the new queue name

Figure 5.16: Enter the name of your queue in the "Name field". Use upper and lower case letters, numbers, the dash character and the underscore character. Queue names may not contain spaces.

Figure 5.17. Setting the input TCP port

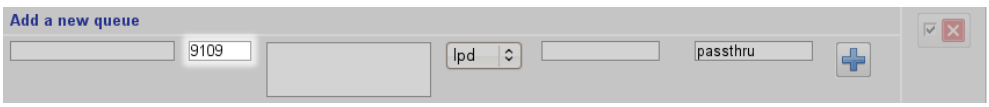

Figure 5.17: Enter a TCP/IP port number for raw port printing. This is the port on which the queue will accept print jobs. A common port for print applications is 9100 which is the default on the PdfPlus. If you will not use raw port printing, simply leave this set to 9100.

Figure 5.18. Setting the queue information

Figure 5.18: Enter a description of the queue here which will help you administer the system.

Figure 5.19. Configuring the output settings

Figure 5.19: Choose an output data transfer protocol and specify the destination address and queue or port. If you are sending data to a printer, the most likely settings here are “LPD”, the IP address or network name of the printer, and “passthru” (usually, a network printer will respond on any LPD queue, so “passthru” is simply a common choice).

Figure 5.20. Adding the new queue

Figure 5.20: Finally, press the “add” button to create the queue.

Figure 5.21. The new queue is ready

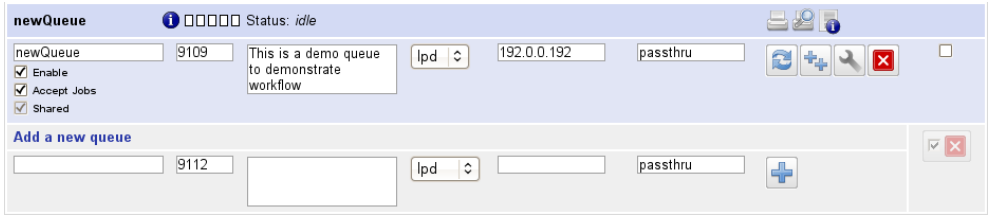

Figure 5.21: The new queue will appear directly above the “Add a new Queue” row. It will be created in the “accepting” and “enabled” states (see the checkboxes to the left of the new queue row), meaning that it will accept new jobs for processing and it will process any accepted jobs immediately.

Figure 5.22. Enter the queue configuration for the new queue

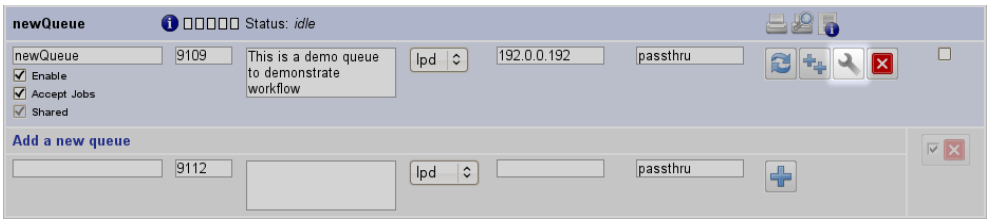

Figure 5.22: Press the “configure” button to enter the queue configuration page for the new queue.

Figure 5.23. Configuring the output PPD driver

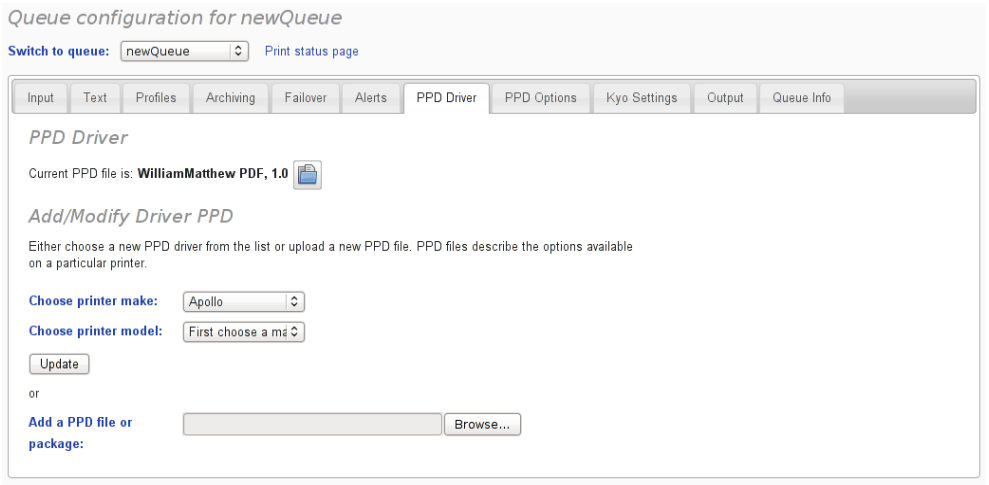

Figure 5.23: The most important configuration setting for the queue is the PPD driver. This is where you define how the PdfPlus should format the document before it is sent to its destination. The PPD driver selects which of the PdfPlus's internal printer driver filters is used to convert the job from the PdfPlus's internal format (PDF) into printer-ready output.

Refer to the section called “PPD driver” if you need to upload your printer manufacturer's own PPD file. Here, we will assume that the PPD file for your printer has already been uploaded to the PdfPlus.

Figure 5.24. Selecting a printer make

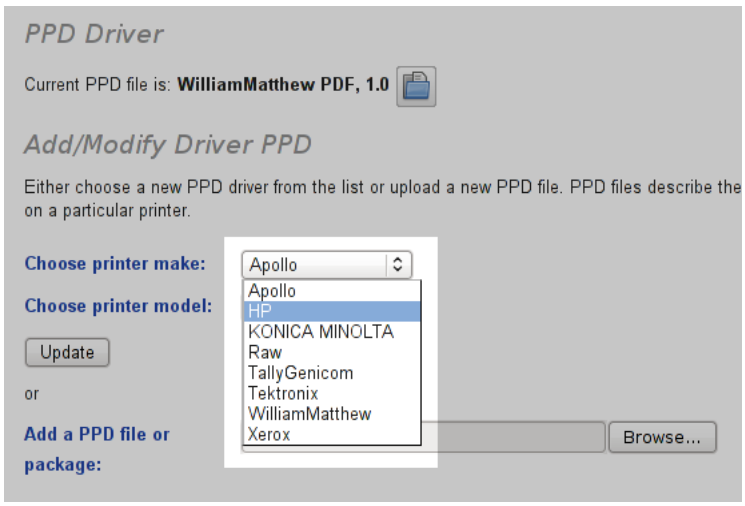

Figure 5.24: Choose the printer make from the drop-down menu. In this example, we're selecting an HP printer.

Figure 5.25. Selecting a printer model

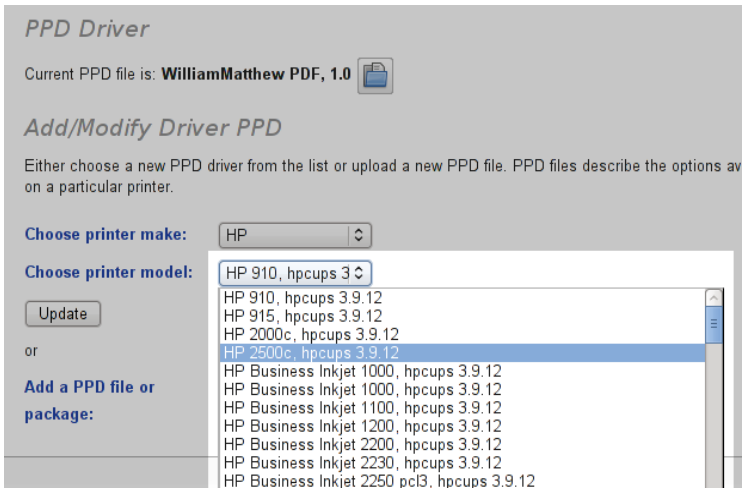

Figure 5.25: Choose the printer model from the second drop-down menu (which will now contain all the HP printers for which which the PdfPlus has an installed PPD file). Here, an HP 2500c is selected.

Figure 5.26. Comitting the changes

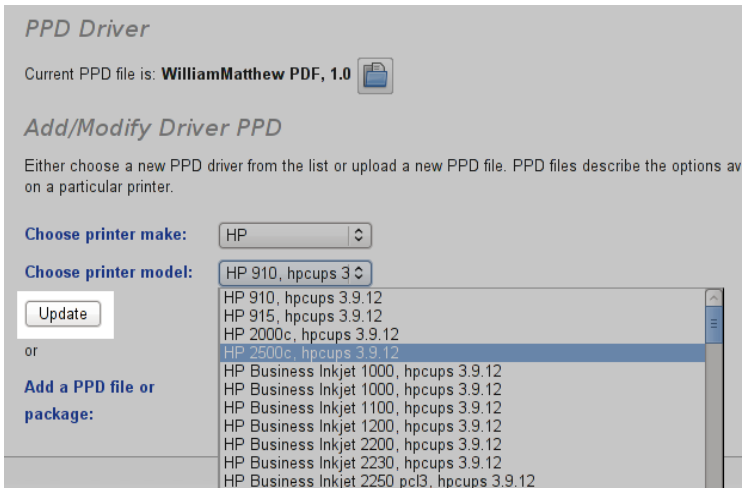

Figure 5.26: Press “Update” to change the the newly selected PPD driver. The “current PPD file” should change to the new value (all new queues have “William Matthew PDF, 1.0” by default).

Figure 5.27. Configuring PPD options

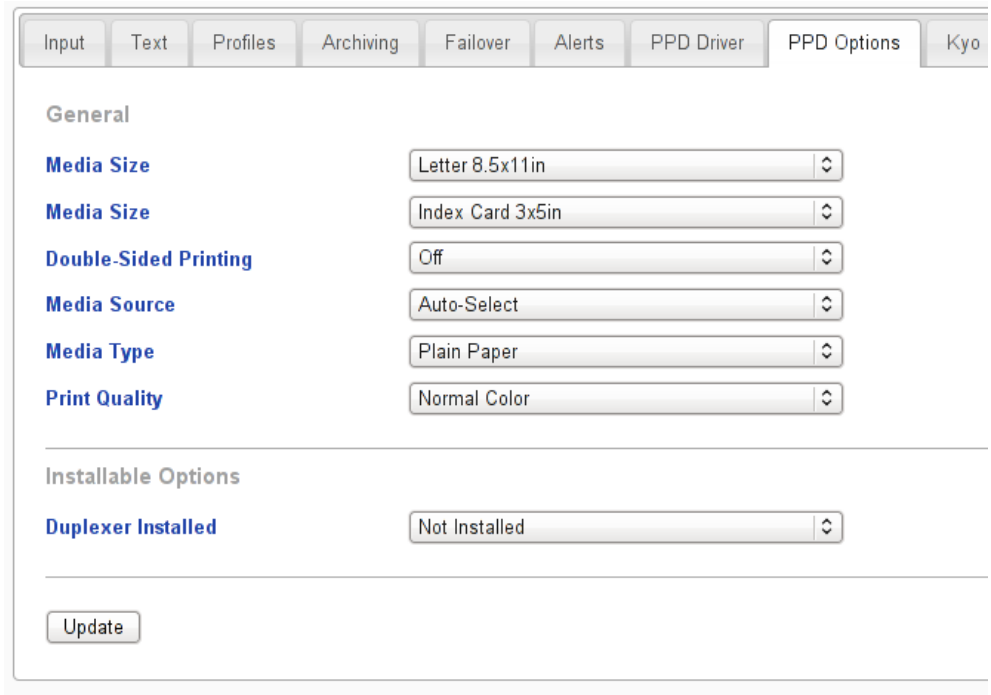

Figure 5.27: Any printer options available for the selected printer will now be available in the “PPD Options” tab. Configure these to your requirements.

The new print queue is now ready for use. You can send data to the queue using LPD, RAW and IPP print protocols. If you need to use a hot folder for transferring print data to the PdfPlus, archive your print data or configure other aspects of the queue, please refer to the sections earlier in this chapter for details.

5.4. Jobs

The **Print Platform > Jobs** page shows lists of currently processing and completed jobs. It shows details for a limited number of jobs; by default the last 500 print jobs are listed. This number can be adjusted in **Print Platform > Platform Settings** by going to the **Cups Control** tab and setting the “MaxJobs” parameter.

Jobs printed on individual queues can be shown using the “Print queues” menu. The size of the job lists can be limited with the “Number of jobs” menu.

If job data is held, in addition to job info (by setting the parameter PreserveJobFiles in the **Edit cupsd.conf** tab of **Print Platform > Platform Settings**), then jobs can be re-printed via this interface. Note that re-printing here will be from the original *input data*, so if any meta-data has been applied to the job previously (such as the current date or time) or if a PDF form was applied and has since changed, the print job may render differently from the originally printed version. For accurate re-printing of jobs, the job archiving and re-print feature should be used (**File Manager > Archive Files**).

5.5. Platform settings

Print settings which are not specific to a single queue are configured via the “Platform Settings” page.

Cups Control

All WML Print Platform products use CUPS for print data transport. This tab allows certain CUPS parameters to be changed by the user, and allows the CUPS daemon to be re-started.

- LogLevel: Determines which log messages CUPS output to the system log.
- MaxJobs: The number of jobs for which CUPS holds information.
- Port: The TCP port on which CUPS listens.

> **Note**
>
> Leave the LogLevel at “notice” if print job reporting features are being used.

Failover

Check “Enable failover mode” if the main print destinations have failed and jobs should be re-routed to the failover queues.

Edit cupsd.conf

This feature is intended for support staff only. This is the main CUPS configuration file.

Edit mime.types

This feature is intended for support staff only. This is the configuration file which determines how the file type of input data files is identified.

Edit mime.convs

This feature is intended for support staff only. This is the configuration file which determines how data files are converted into the PdfPlus's internal format, PDF.

Chapter 6. Template Management

6.1. Introduction

The PdfPlus's main tool for data manipulation and document creation is the *template manager* (Figure 6.1). This WYSIWYG tool allows you to build up a *template* which can be applied to print jobs as they are processed by the PdfPlus.

A *template* is a collection of *overlay elements* which are individual items of text or images. The contents of the template are merged with the incoming print data to create the finished print document.

A unique feature of the WML PdfPlus is that all data manipulation is applied natively in the PDF format. This means that the output document is always available as a PDF for archiving or electronic delivery.

Figure 6.1. The brower-based template manager

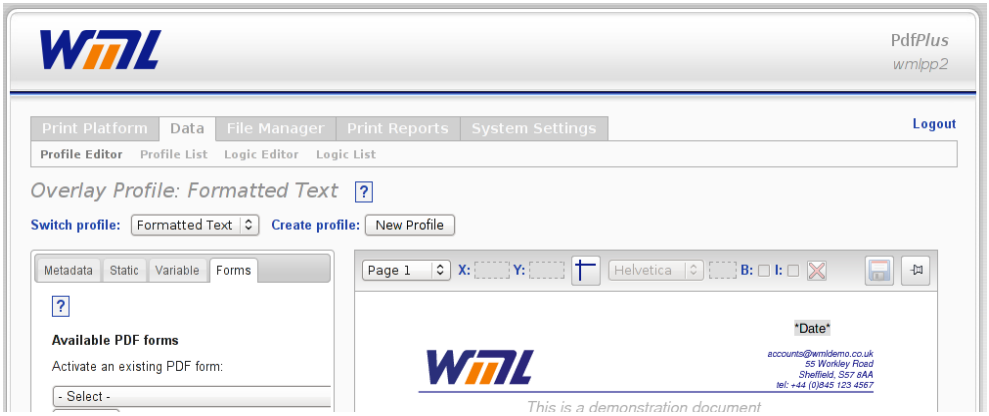

It's easy to build up a template using the PdfPlus's template manager. Simply upload your PDF Forms and drag and drop text elements into position. It's a quick, easy job to build up a template which will enhance the quality of your print output.

Figure 6.2. A finished example template

Figure 6.2 shows a finished template containing date and time realtime data elements, a static element, and the same PDF form as shown in Figure 6.5.

Figure 6.3. The result of the example template viewed in the Adobe PDF Reader

Figure 6.3 shows the result of passing a simple input print job to be processed by the example template.

6.2. Overlay elements

An **overlay element** is an item of data to be applied to a page of a print job. This might be a piece of text, a group of texts, an image or even a complete PDF. In fact, these are the four categories of overlay elements; **text elements**, **data group elements**, **image elements** and **PDF form elements**.

Each overlay element has an **element type**. The different types are:

- Realtime data: A text element whose value is generated as the data is processed (e.g. the date).
- Static data: A text element which always has the same value.
- Variable data: A text element which is defined with a key. The value is provided at processing time with input print data. Variable data elements can be simple, single elements (e.g. "customer_name" filled with the value "Mr Smith") as well as **Data Groups** to lay out tables and multiple paragraphs.
- Data group: A set of text elements which should be laid out over a given area. Used for documents with variable size, such as those with a variable number of paragraphs, or variable length tables.
- PDF form: A PDF file containing one or more pages.
- Barcode: An image element, generated by converting a piece of text within the input print data into a barcode.

An element can be added to a single page of a document (perhaps page 1 or page 3) or it can be added to all pages of a document.

Figure 6.4. Three text elements; a date, a time and a piece of static text

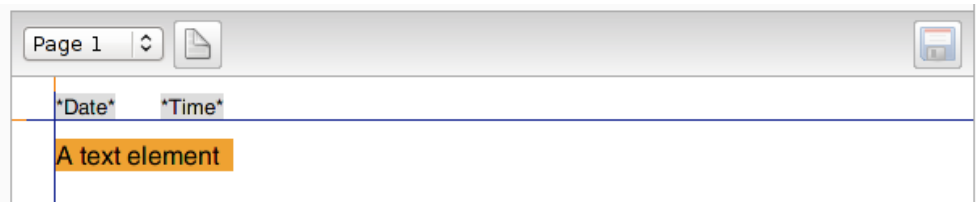

Text elements have position and font specifications, controlled by the "Element editor" (seen in Figure 6.1 and Figure 6.7). Image elements are simply given a position. PDF form elements are laid onto the page, as if they were a transparency on a projector.

Figure 6.4 shows three text elements, positioned towards the top left of page 1 of the template. The snap-to lines are visible and the static text element is highlighted in orange, showing that it is the currently selected overlay element.

Figure 6.5 shows a single PDF form element containing a WML logo, which is active for page 1 of this template.

PDF form elements may have multiple pages. If you add a multi-page PDF form to page 1 of a template, then page 2 of the form will print on page 2 of the template, page 3 of the form on page 3 of the template and so on, as long as there are enough pages in the template or incoming print job.

Figure 6.5. The logo seen here is a PDF form element, applied from an uploaded PDF file

6.3. Overlay templates

An *overlay template* combines text elements, image elements and PDF forms which can be applied to individual pages, or to every page in the input print job.

It also includes a number of *template settings* such as the template name, the number of pages for the template and the page size.

An overlay template is stored in a single, XML-based configuration file on the PdfPlus. Overlay templates can be created on one PdfPlus and then exported and copied onto other PdfPluss, either individually using the Template Manager, or as a group using the facilities for backing up the PdfPlus configuration in **System > Settings**.

An overlay template can be "bound" to a queue, so that all data passing through that queue will be processed by the overlay template. The overlay template is bound to the queue in the queue configuration, see Chapter 5.

Alternatively, the WML print filters can select an overlay template based on the existence of text triggers within the input print job. For example, an overlay template named "CompanyInvoice" might be triggered when a document containing the text "INVOICE" were processed. For more information about logic-based triggering, see Chapter 7.

6.4. Overlay template list

To create the first overlay template, or to view a list of templates, select the menu item **Data > Template List**.

Figure 6.6 shows an example of the template list. It is a table with one line for each overlay template. The template name is shown, as well as the *document basis*, the number of overlay elements in the template and a column of action buttons.

> **Note**
> The meaning of *document basis* will be described later.

Figure 6.6. The overlay template list

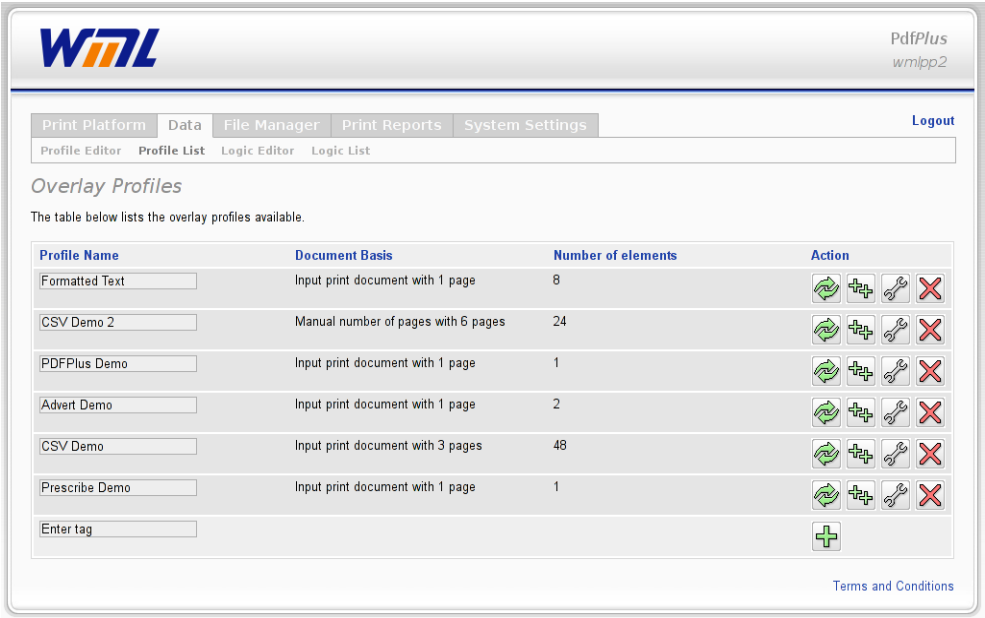

To create a new overlay template, enter a name for the template. The name should be unique because it is used as an identifier for selecting the active overlay template (either manually or by means of a trigger in the input data). Click on the add button ( ) to create the new template.

You may also duplicate an existing overlay template, by clicking on the duplicate button ( ) for the template that you wish to copy. A new template will be created and added to the list.

While a newly created overlay template does not contain any overlay elements, a template created by duplication will contain copies of the elements in the duplicated template.

You can change the name of an existing overlay template by editing the value in the text field and clicking on the update button ( ).

> **Note**
> If you modify the name of a template which has been assigned to a particular queue, you will need to re-assign the template to the queue using the new name.

From the template list, it is possible to “drill down” into the template manager for any of the listed templates. Edit a template by clicking on its configure button ( ).

6.5. Overlay template manager

The overlay template manager was shown previously in Figure 6.1 and is presented in full view below (Figure 6.7).

Figure 6.7. The overlay template manager

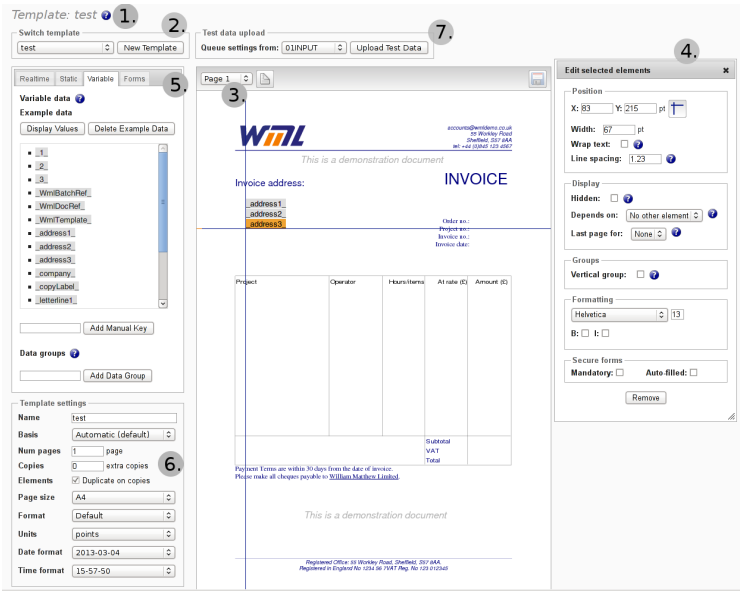

The template manager consists of the following components:

1. The title - this tells you the name of the current overlay template.
2. A top toolbar which allows you to switch to a different template, or create a new one.
3. The template “canvas”, showing you how the overlay elements will print. At the top of the canvas is the “canvas toolbar”.
4. The element editor (to the right of the canvas) which contains settings for the currently selected overlay element(s).
5. The overlay elements box, from which you can drag text elements onto the canvas, and activate PDF form elements.
6. The template settings box, which allows you to change the overlay template's attributes.
7. The Test Data box, which contains an upload field which you can use to upload an example print job. This is used to test your overlay template without the need to fully set up a print queue.
8. The template download box (not visible in the image; below the Template settings box). Allows the user to save an archive of the current template.

Saving changes to the template

The template manager is a browser-based front-end to the PdfPlus appliance. It sends messages to a back-end program on the PdfPlus which writes changes to a configuration file.

In general, changes made in the left column are immediately saved on the PdfPlus without any further action. In contrast, changes to the position, width, font size and weight of any text elements are *not* immediately saved. Any such changes will cause a "modification warning" to display at the top of the canvas. Figure 6.8 shows an example of this warning.

To commit changes to the PdfPlus it is necessary to press the save button ( ).

Figure 6.8. This template has been modified

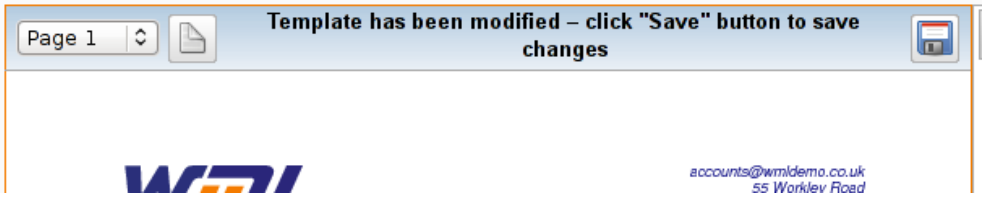

Template management

Switch template

Immediately below the title are two controls: A menu which allows you to switch to a different template and a button to create a new template.

If you try to switch to a different template and you have made changes to text elements which have not yet been saved, the PdfPlus will open a dialog to ask if you are sure you wish to switch and lose your un-saved changes.

If you click the "Create template" button, a dialog will open asking you for a name for the new template (Figure 6.9). You can also choose to create a new template by uploading a template archive (as saved under **Template operations**).

Figure 6.9. The new template dialog box

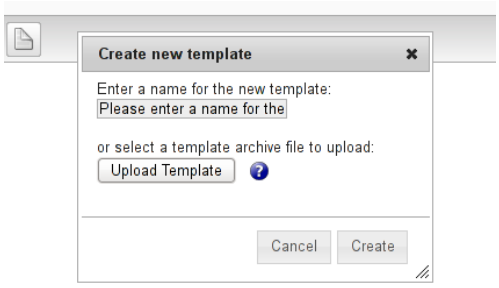

Template operations

Below the **Template settings** box is the **Template operations** box. This allows you to duplicate the template or save the template to a file.

Pressing **Duplicate** will create a copy of the current template with an automatically generated name (the name of the original template with "copy" appended).

To save the template, press **Create Archive**. This makes up an archive file containing the layout of the template and all PDF forms which are part of it, and allows you to download it with your brower. The file is called **template.tar**.

> **Note**
>
> To upload the template.tar file to another PdfPlus, so that the template can be used on that system, use the **New Template** button in the **Switch template** box. The **Create new template** pop-up has an **Upload Template** button.

Template settings

The template settings box provides a number of settings which are "template-wide". That is, they affect all elements in the template, or they affect the way the template operates. Figure 6.10 shows the template settings box.

> **Note**
>
> Any change made in any of the fields in the template settings box are *immediately* saved on the PdfPlus.

Figure 6.10. The template settings box

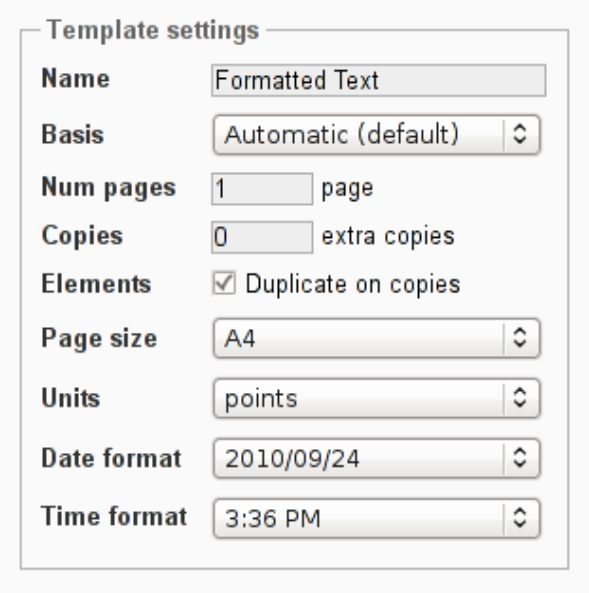

Name

This is simply the name of the template. The template name is modified by changing the value here. (It is also possible to change the name in the template list page, see Section 6.4).

The template name may contain letters, numbers, spaces, dashes and underscores.

> **Note**
>
> If you modify the name of a template which has been assigned to a particular queue, you will need to re-assign the template to the queue using the new name.

Basis

The *Document Basis* determines the way in which the input print document and the template, with any PDF forms, are merged together into a finished document.

The PDF document format has a feature which allows new objects (such as lines of text or images) to be added to the pages of the document simply by appending data to the end of the file. The PdfPlus makes use of this feature.

The PDF to which all other objects are appended is the “document basis”.

> **Note**
>
> The number of pages in the completed document will be the same as in the document basis.

The Basis is set to one of four possible values, which are described below.

Basis: Input print document

If the input print document is used as the basis, then the output print jobs will always have the same number of pages as the input print document.

This mode is used when you expect print jobs of different sizes to print on your queue.

> **Note**
>
> Use input print document mode if you want to apply the same template elements to the first few pages of every document, but your print jobs are of variable length (in pages).

Basis: Manual number of pages

If you specify "Basis: Manual number of pages", then the PdfPlus will create an empty PDF containing the specified number of pages and use that as the basis for building the output document.

In this mode, output print documents will *always* contain the number of pages specified in the "Num pages" field in the template settings box.

> **Note**
>
> Use this mode if your input print documents contain fewer pages than you would like to print.
>
> For example, you could add terms and conditions to the second page of your single-page invoice documents by setting Basis: Manual number of pages; Num pages: 2 and placing a PDF form containing the terms on to page 2 of your template.

Basis: Automatic

This is the default basis for new templates.

In "Basis: Automatic mode", the WML print filters will choose between "Basis: Input print document" mode and "Basis: Manual number of pages" mode. It makes this decision based on the number of pages in the input print document, and the number of pages in the template.

If the input print document contains the highest, or equal highest number of pages, the input print document is used as the basis for building the output document. Otherwise, the filter creates an empty PDF containing the number of pages specified for the template and uses that as basis.

Basis: PDF form

In certain cases, it is desirable to use one of the PDF Forms in the template as the document basis.

Each active PDF form *on page 1* of your template will appear as an option in the Basis menu Like this:

```
PDF form: myfile.pdf
```

For example, you might add a multi-page PDF to your template and use this as the Basis. The advantage is efficiency. It is less processor-intensive to use the largest PDF form as the basis when compared with the "Basis: Manual number of pages" mode.

> **Note**
>
> When you are using the PdfPlus to print structured input data (CSV or XML formatted key-value pairs) it is highly recommended that you use a PDF form as your document basis.

Num pages

The number of pages in the template. The exact meaning of this number is dependent on the document basis.

In "Basis: Input print data" and "Basis: Automatic" modes, this is *the number of pages which may be modified by the template*. If Num pages is set to "2", and a single page document is printed, the elements on the second page of the template will be ignored.

In "Basis: Manual number of pages" mode, Num pages is *the number of pages which the output document will contain*. The output document will always have this number of pages, even if some of the pages are blank.

If a PDF form is chosen as the document basis, then this field will be greyed out and set to the number of pages in the chosen PDF form.

Page size

This is the expected page size of the input print data. You should match this to the size of the documents which you will be printing.

If you do not match the page size of the input print data and the page size specified here in the template, your overlay elements will not print in the correct positions.

> **Note**
>
> If the WML print filters detect that you printed a A4 document through template set to US Letter (or vice versa), it will print a warning page to the effect.

Units

The units in which all overlay element positions will be specified. This defaults to "points" (1/72 of an inch) which is commonly used internally within PDF documents. Other options are *mm*, *cm* and *inches*.

Date format

This is a drop-down menu with a number of date formats such as "24th September 2010" and "2010-09-24".

The date format will be applied to all Date realtime data elements on the template.

> **Note**
>
> In the current revision of PdfPlus, it is not possible to have differing date formats within the same template. For example, you could not have Date realtime data element on page one filled in as “14th Sept 2010” and one on page two filled in as “2010-09-14”.

Time format

Like the date format menu, this is a drop-down menu with a number of formats for printing the current time in Time realtime data elements.

> **Note**
>
> In the current revision of PdfPlus, it is not possible to have differing time formats within the same template.

The canvas

Figure 6.11. The template manager's “canvas”

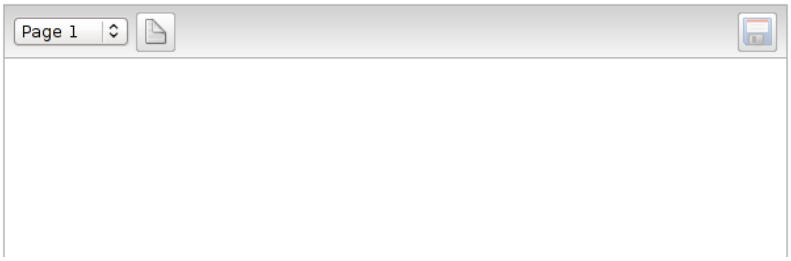

The canvas is the main work area of the template manager (see Figure 6.11 which shows an empty canvas). It displays the current template page as it would print. Any active PDF form elements are displayed, along with text elements and “snap-to” lines, if they are enabled.

You can select one or more elements (they will turn orange when selected) and move them by dragging. Controls in the element editor allow changes to the font as well as enabling you to fine-tune the position and width of the fields.

The canvas toolbar

At the top of the canvas is the *canvas toolbar*. This contains a menu to switch to other pages of the template (the *pages menu*), a media control button and a save button to write changes to the PdfPlus.

The pages menu

In Figure 6.11, the canvas of page 1 of the template is displayed. This is seen in the menu at the far left of the canvas toolbar (the *pages menu*). To modify overlay elements on a dif-

ferent page, select the page from the pages menu. The pages menu contains one entry for each page in the template and additional entries for several special pages: “Others”, “Odd”, “Even” and “All pages”.

> **Note**
>
> New pages are added to the template in the template settings box (see “Num pages”).

“Others”

By selecting “Others”, you are placing overlay elements *which should appear on any page of an input document which is not covered by the template*.

For example, Imagine your template contains a single page, on which you place a company letterhead. Now imagine you place a simpler PDF form containing only a company logo on the “Others” page. If you print a one page document through the template, the letterhead will be printed on page one. If you print a two page document, the letterhead will print on page one and the logo will print on page two. If you print a one hundred page document, the letterhead will be printed on page one, and on every other page, the logo will be printed.

“Odd/Even”

By selecting “Odd” (“Even”), you are placing overlay elements which should appear on any odd (even) page of the print job.

“All pages”

If you select “All pages” from the pages menu, then any elements you place will appear on every single page printed through the template.

> **Note**
>
> Text elements which are placed on the “All pages” canvas will be shown in blue text. These elements will appear on all other pages of the template. For example, if you switch to page 1 of the template, the elements you placed on the “All pages” canvas will appear alongside those which will only appear on page 1.

Saving changes

To commit changes press the save button ( ).

The element editor

The element editor contains fields controlling the position, width, font-face, font-weight and font-size of selected elements.

Position and width

The *X* and *Y* fields show the current X and Y position in the current units, with respect to the top-left of the page. These values can be edited and the selected element will move to the new position.

> **Note**
>
> The units default to “points”. The units can be changed in the template settings (see “Units”).

The *Width* field shows the element width in PDF units (points). The width is the required width of the element when printed. It is used in the WML Secure Forms module to size the editable field for each variable data text element. You can also check the *Wrap text* option to wrap the element text to the specified width. The *Line spacing* field governs the spacing within a word-wrapped text element. The line spacing is specified as a proportion of the font-size.

Formatting

The font face can be selected from the drop down menu. Directly to the right of the font menu is a field to specify the font size.

The available fonts are: Times New Roman, Helvetica, Courier, Symbol and Zapf Dingbats. These are the basic “built-in” fonts which are available in every PDF document.

The font drop down menu also provides options for barcoding. If enabled, you can also choose an option to convert a text element to a barcode image.

The current font can be emboldened or italicized by checking either the *B* or the *I* boxes, respectively.

Deleting elements

To delete selected elements, press the “Remove” button, or press the delete key on your keyboard. As with other changes to elements, the changes will not be saved on the PdfPlus until the save button is pressed.

Snap-to lines

The template manager provides “snap-to” guidelines which can be enabled and disabled as required. To enable the snap-to lines, press the snap-to lines button ( ). The snap-to lines can be seen in Figure 6.4.

If you drag an element towards the snap-to lines, the element will dock with the lines, allowing you to line up a set of elements.

To disable the snap-to lines, press the snap-to lines button again.

Working with multiple elements

It's possible to select multiple elements and apply the same font-face and font-size to the set. Multiple selected elements can also be dragged as a group.

You can also modify the X and Y positions for a group; however, if you select multiple elements, and they are offset from each other, then changing the position will maintain the offset between them. The position of all selected elements will change by the amount that the first element moves.

The overlay elements user interface

The overlay elements user interface box, shown in Figure 6.12, is the “source” for all overlay elements. From here you can manage realtime data, static data, variable data and PDF forms. Text elements are dragged onto the canvas. PDF forms are “activated” for the current page of the template.

> **Note**
>
> For a description of each overlay element type, please see Section 6.2.

The overlay elements interface consists of four tabs, one for each type of overlay element. Each tab is described here.

Realtime elements

Figure 6.12 shows the realtime data tab. Four realtime data elements are displayed: Date; Time; Title and Source IP. The value of these fields is filled in when the job is processed (i.e. just as it is printed).

Figure 6.12. Overlay elements - Realtime tab

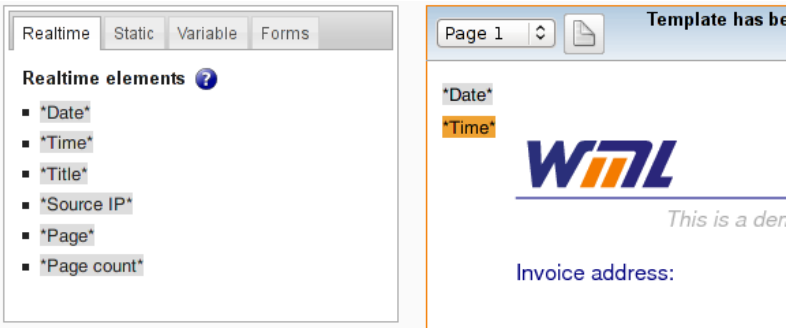

To place any of these elements on the canvas, drag the element from the overlay elements realtime tab and drop them in the desired location.

Figure 6.12 shows a date and a time element placed on the canvas (the time element is selected and is highlighted orange). Note that realtime data elements are enclosed by '*' characters. The date element appears on the canvas as “*Date*” the Time element as “*Time*”.

The date and time realtime data elements are intended to be used in time-stamps to prove the date at which a job was processed for printing or archival. The format in which the date and time are printed is configured in the template settings.

The title realtime data element is filled in with the job title. This may be set by the user if the print data has been transferred to the PdfPlus using the LPD or IPP protocols, which allow the specification of a job title. If a hotfolder has been used to transfer the print data onto the PdfPlus then the filename of the print file in the hotfolder will be used. If raw TCP/IP data transfer has been used, or if the title has not been set when using LPD or IPP; the title will print as “unknown”.

The “Source IP” realtime data element will contain the IP address of the sending computer system if either the LPD or the raw TCP/IP protocols have been used to transfer print data onto the PdfPlus.

> **Note**
>
> Unfortunately, the sender's source IP address cannot be determined when the IPP protocol or hotfolders have been used for print data transfer.

Static data elements

Static data is the simplest overlay element. It is an unchanging piece of text (a word or phrase) which always prints in the same way. Static data can be used to make simple amendments to documents without the need to go back to the original artwork editor.

Figure 6.13. Overlay elements - Static Data tab

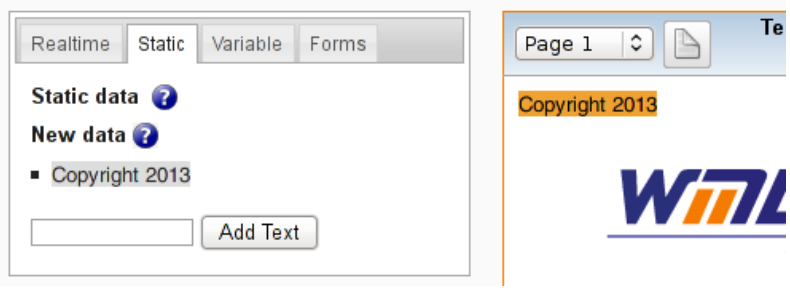

Figure 6.13 shows the static data tab, with a single static data element containing the text “Copyright 2013”. To place the element, drag and drop it in the usual way. Multiple copies of the static element may be placed on the canvas.

To create a new static data element, with different contents, fill in the blank text field and press the button “Add Text”. The new element will appear below “Copyright 2013”. You can now drag this element onto the canvas.

> **Note**
>
> When a template is accessed, the static data tab will show only static data elements which are placed on the template. If you add a static element with the “Add Text” button, then don't use it, it will not be available the next time you load the template in the template manager.

Variable data elements

A variable data element is a text element with a “key”. When CSV or XML input print data is received, specifying the value of this key, the value is printed at the location specified by the element and in the specified font.

The use of variable data input - also referred to as “structured data” - is covered in greater length in Section 6.7: “Using structured input data”.

Figure 6.14. Overlay elements - Variable Data tab, initial empty view

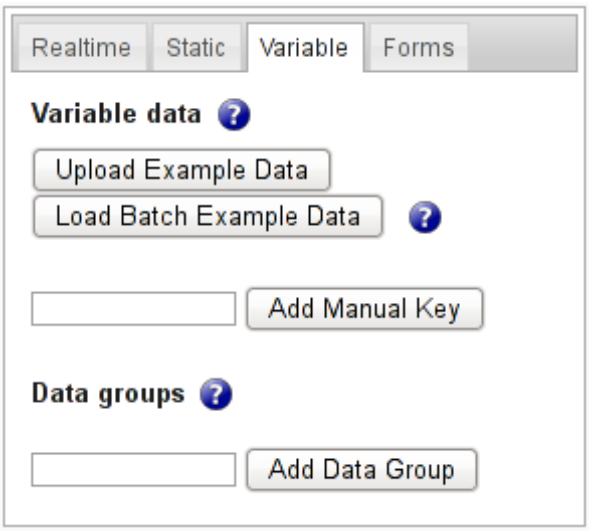

Figure 6.14 shows the variable data tab in its initial, empty state. There is a button to "Upload Example Data" and a field to add a manual key. These are the two methods you can use to create a variable data element. First, we'll cover manually specified keys.

Figure 6.15. Overlay elements - Variable Data tab after adding a manual key

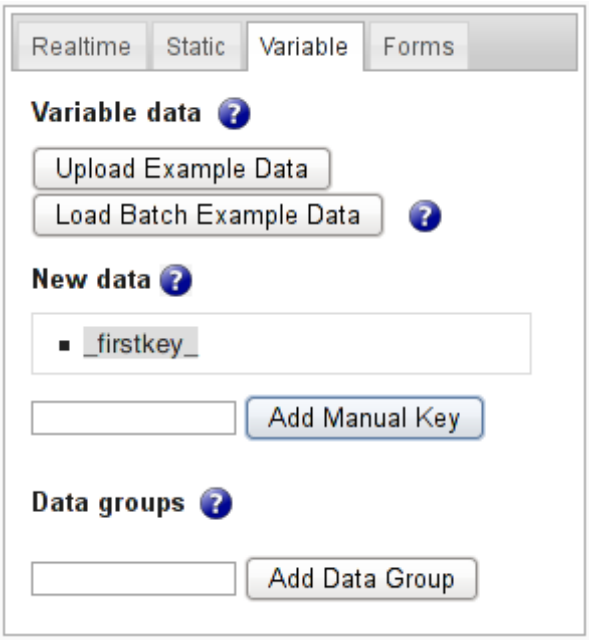

Figure 6.15 shows the result of entering “firstkey” into the manual key text field and pressing “Add Manual Key”. A heading “New data” has appeared with a single item “_firstkey_”. Note the underscore characters enclosing the key. These indicate that this is a variable data element and are not part of the key. They distinguish variable data elements from realtime data elements, which are enclosed by '*' characters and from data group elements, enclosed by '/' characters.

Figure 6.16. Overlay elements - Variable Data tab showing a saved key

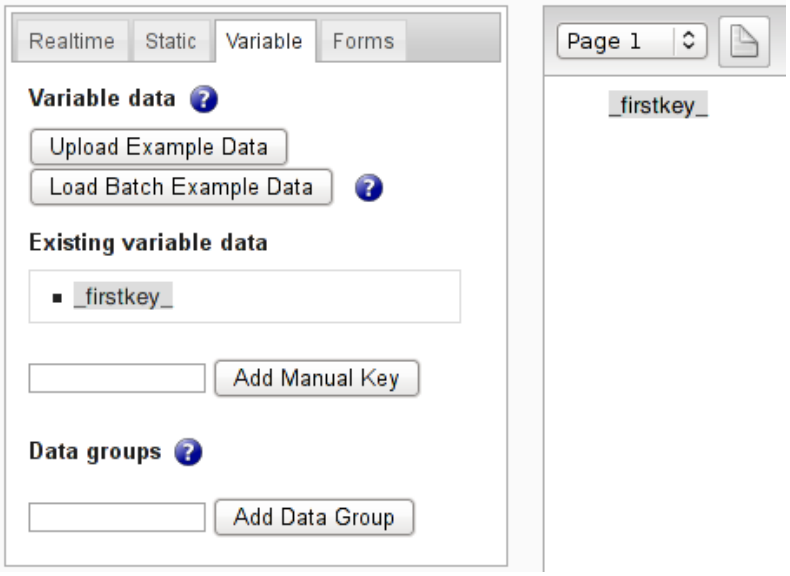

This variable data element has not yet been placed on the template canvas. Figure 6.16 shows the variable data tab after the key has been placed. “firstkey” is now in a (short) list under the heading “Existing variable data”.

Manually specifying keys is not the most convenient way to add variable data elements to an overlay template. A more efficient method is to upload an *example data file*. This is a structured data file, arranged in CSV format or WML's simple key-value XML format. The template manager processes this file and extracts the keys, and the example values from the file.

Example 6.1. An example variable data file in CSV format

```
name,Mr Jones
address,24 Tenants Drive
```

Example 6.1 shows a very simple example data file containing two key-value pairs. This file is formatted as comma separated values. The first column contains the keys - in this case “name” and “address” - and the second column contains the corresponding values - “Mr Jones” and “24 Tenants Drive”.

> **Note**
>
> The example variable data file should be exactly the same as the file which would be sent to print. It is simply an example structured data print file.

Figure 6.17. Overlay elements - Variable Data tab showing keys

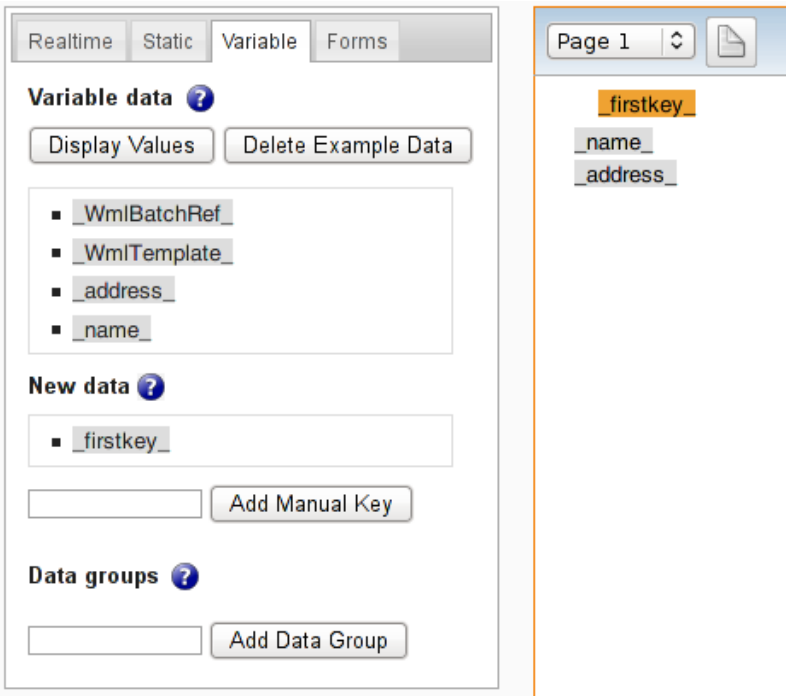

Figure 6.17 shows the result of uploading the example shown in Example 6.1. The “Upload Example Data” button has been replaced by two new buttons and a list of the keys in the uploaded example data (“address” and “name”).

Figure 6.17 also shows firstkey, name and address placed on the template canvas. In common with all other text elements, multiple copies of each variable data element can be placed on the canvas. Here, the “name” key could be printed in several different locations on the same page.

The new buttons are “Display Values” and “Delete Example Data”. “Delete Example Data” simply clears the example data from the template, leaving only any existing variable data elements (i.e. those which have been placed on the canvas and saved).

Figure 6.18. Overlay elements - Variable Data tab showing values

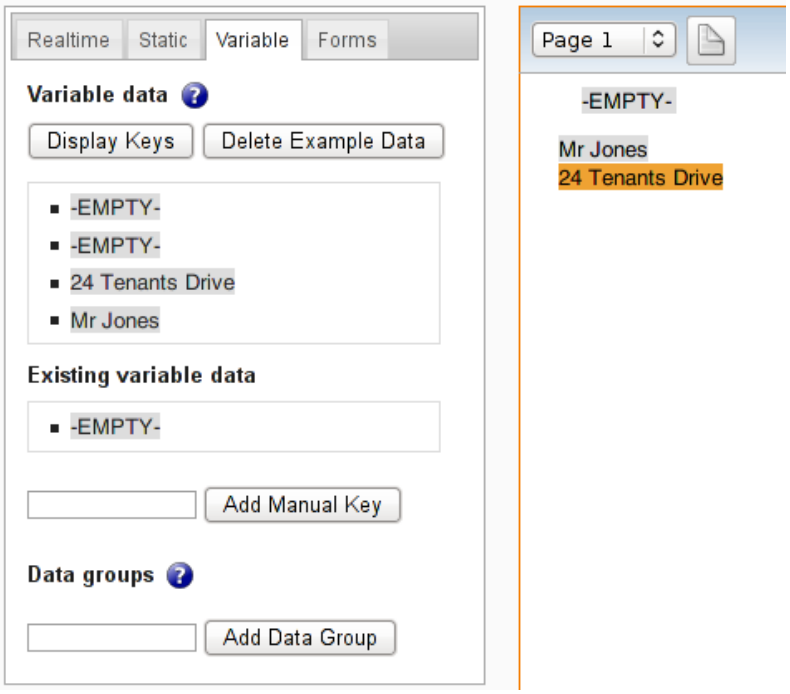

"Display Values" will change the view in the template canvas to show the example value of each variable data element. This is shown in Figure 6.18. Showing the example values provides a better impression of the way the document will print.

> **Note**
>
> The manually created key "firstkey" has no example value and so it is displayed as "-EMPTY-".

Data groups

Variable data may be brought together in *data groups*. A data group is an element which allows you to control the display of a sequence of related data, with the placement of each data member being determined automatically.

When the input data is processed, the elements assigned to a data group are positioned in sequence on the page, according to the size and spacing attributes given for the data group. If the elements fill up the space available for the group then a new page is created, and subsequent elements are placed there. This allows the flexible layout of variable length documents such as letters and invoices.

Figure 6.19. Overlay elements - Adding a data group

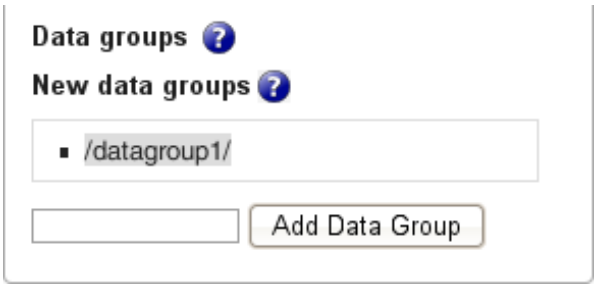

Figure 6.19 shows the result of entering “datagroup1” into the data group text field and pressing "Add Data Group". A heading “New data groups” has appeared with a single item “/datagroup1/”. Note the slash characters enclosing the data group name. These indicate that this is a data group element and are not part of the data group name.

When the data group is placed on the canvas, its name will move to a list under the heading “Existing data groups” (Figure 6.20). Data groups are highlighted in pink on the canvas, with a grid background that indicates the area used by the group. The area can be adjusted using the width and height fields in the element editor. Other attributes, such as font size and line-spacing, can be set in order to control how the elements in the data group will be formatted.

Figure 6.20. Overlay elements - Canvas with data group

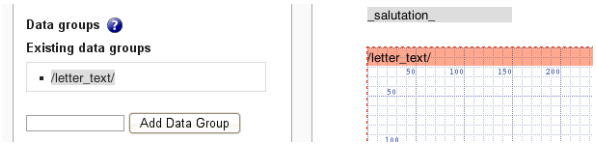

The first data group added to the canvas is the “master” data group. You will be able to specify a set of variable data elements to be included in the group, and to control the layout of the group. Subsequent copies of the data group will define “continuation” areas, into which elements defined for the master group may overflow. These can be on the same page or on a special *continuation page*.

> **Note**
>
> It is only possible to edit the variable data keys for a group through the master group element. You can adjust the layout and formatting of a continuation group, but the list of variable data will be disabled.

Some examples of using data groups to display variable input data can be found in Section 6.8.

Forms

Figure 6.21 shows the PDF forms tab. This interface has a menu to select and activate PDF forms which have already been uploaded to the PdfPlus, a button to upload new PDF forms to the appliance and a list of active forms. In this example, no forms are active for the template.

Figure 6.21. Overlay elements - Forms tab with no active forms

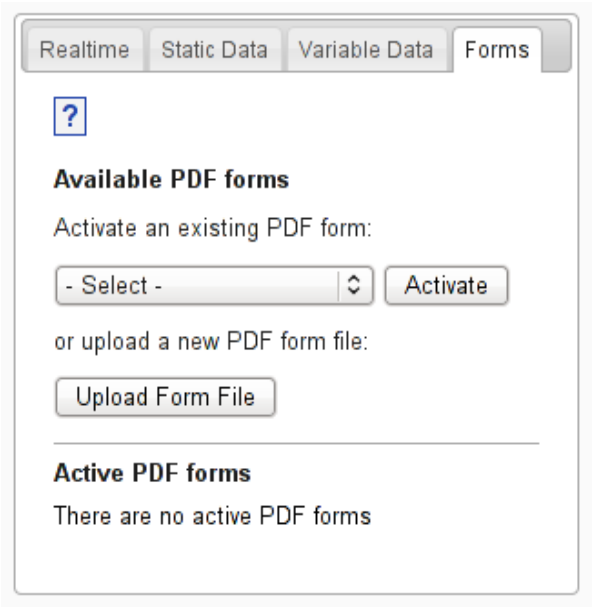

Any PDF forms which you wish to use with the PdfPlus must first be uploaded to the appliance. Click on “Upload Form File” and browse to your PDF. The file will be uploaded to the PdfPlus and processed by the internal software to ensure it can be successfully parsed.

> **Note**
>
> If the PDF cannot be parsed, perhaps because it contains a feature unsupported by the WML PDF software, then an error message will display. Please contact your vendor with details of the problem, ideally sending them a copy of the problem PDF.

Figure 6.22. Overlay elements - selecting a PDF form to activate

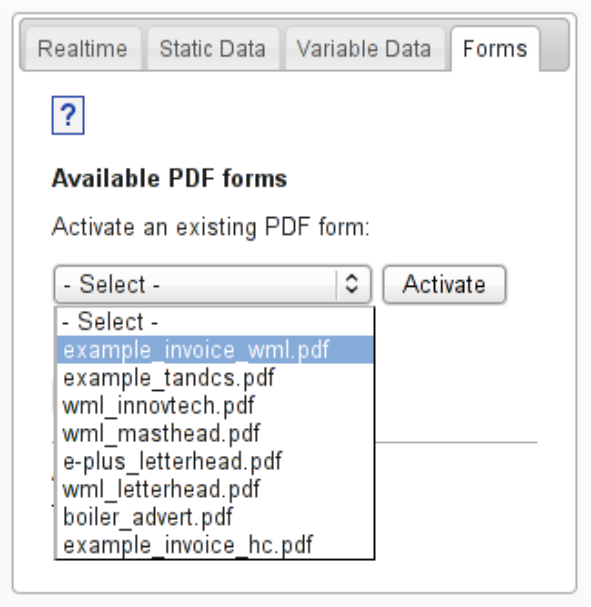

Once you have uploaded your PDF forms, select a form to activate. Choose the form in the "Activate an existing form" menu and press the "Activate" button. An image of the first page of the PDF form will appear on the canvas and box will appear in the "Active PDF forms" list. This will appear as shown in Figure 6.23.

Figure 6.23. Overlay elements - Forms tab with one active form

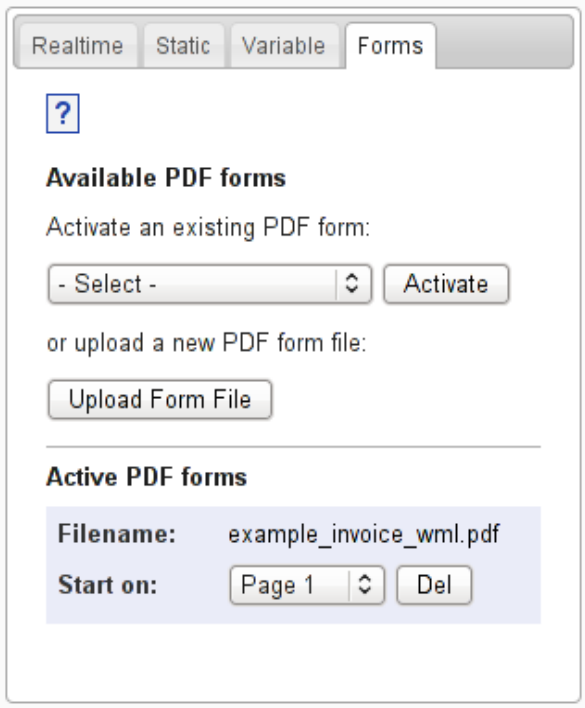

The PDF form will be activated on the currently selected page of the template. The first page of the form will appear on the currently selected template page. If the PDF form should appear on all pages, then change to the “All pages” template page, and then activate the form there.

It's possible to move an active PDF form from one page to another. Each form has a menu listing all the pages in the template. Simply choose a new page for the form (Figure 6.24). To delete the form, press the “Del” button and it will fade out.

> **Note**
>
> All the actions relating to PDF forms are immediately saved on the PdfPlus.

Figure 6.24. Overlay elements - Changing the start page for a PDF form

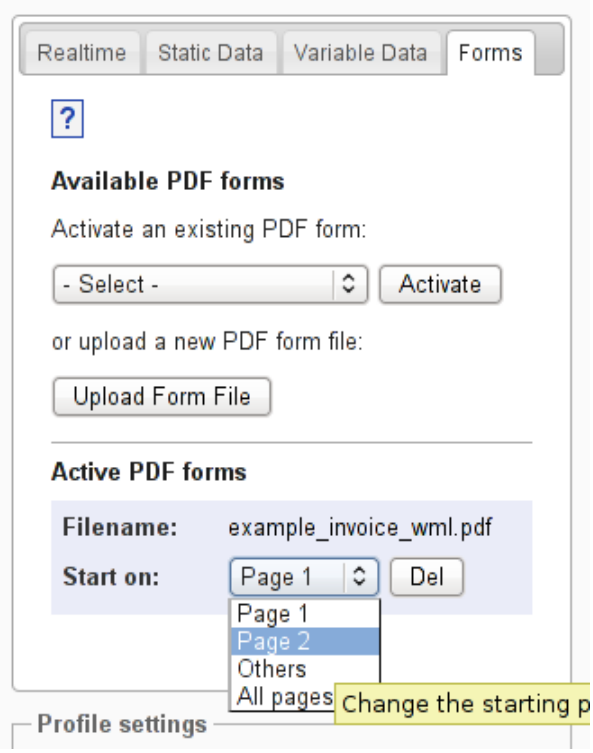

Testing print data

Once you have created your template you will want to verify that it produces good output. The **Test data upload** interface provides the facility to upload an example print file via the web interace. The test data file is processed as if it were a print job and a PDF is returned to your browser. This allows you to test your template without the need to set up a print environment.

Figure 6.25. Test data upload box

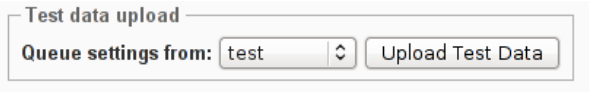

Figure 6.25 shows the test data upload box. It consists of a menu to choose the queue whose settings should be applied to the test, and an upload button.

You can upload any data which can be processed by the PdfPlus. The print data format will be automatically identified and the data will be passed through the relevant WML data filters to apply the template. If you have a browser plugin to view PDF files, the resulting file will appear in a new browser tab or window.

> **Note**
>
> If your browser blocks pop-up windows, then you will need to allow pop-ups for the PdfPlus.

It's necessary to choose a queue because some print jobs can only be processed with the help of per-queue settings. An example is CSV or text data. There's no easy way for a machine to distinguish between comma separated values and plain, formatted text, so one of the per-queue settings is an option to tell the queue to treat text data as CSV or formatted text. If you are testing CSV print data, then ensure you test with a queue which is set to treat text data as CSV.

Other per-queue settings which may affect the output from your template include Prescribe emulation settings and formatted text emulation settings.

> **Warning**
>
> Some care should be taken when choosing the **Queue settings from** option. If you choose a queue which applies document content control (DCC), you may see a status PDF from the DCC system, instead of the expected PDF generated from your template. This is because DCC typically outputs onto a separate **DCC output queue**.

If the print file is accepted for processing, but there is an error in your template, or a problem with the queue you are using for the test, you may see a WML error page PDF appear in your browser. This will be a plain document stating that an error occurred while manipulating the PDF. The message in the error may indicate a solution to the problem. If not, please contact your vendor for assistance, with a description of the error.

6.6. Overlay Templates: Example workflow

This section steps you through the basic work-flow of creating and applying a template. This will not cover every feature of overlay templates, but provides a starting point for a new user and an overview of the process.

Create a print queue

If you do not already have a print queue configured, follow the instructions in Section 5.3 (Chapter 5) to create a queue. In this workflow example, we'll assume it's called “newQueue” (this comes into play later on when we test the template).

Create PDF forms

You can use almost any “artwork” program you like to generate your PDF forms. You could choose a word processor such as OpenOffice.org or Microsoft Office or a vector graphics program such as Adobe Illustrator or the open source Inkscape. If the program does not have its own PDF output function, then a PDF printer driver can be used to generate PDF output by “printing to a file”.

Figure 6.26. Printing to PDF with OpenOffice.org and PdfCreator

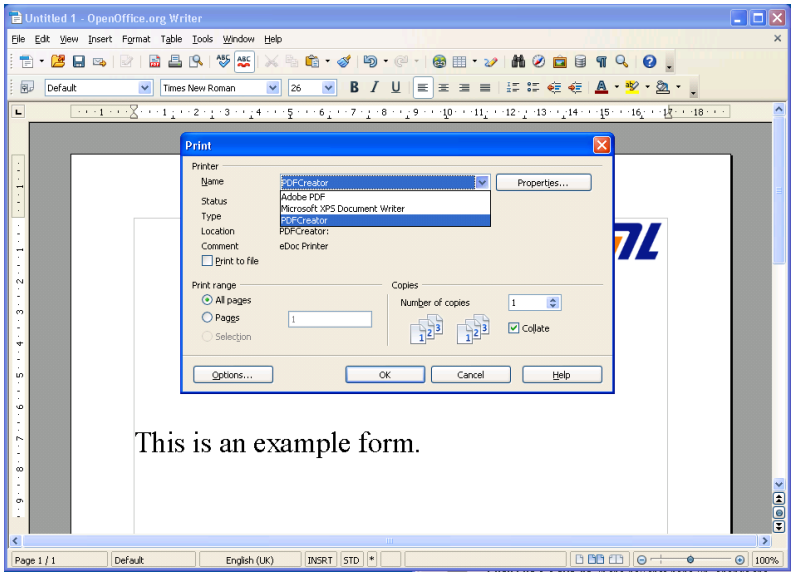

Examples of PDF printer drivers which can be used include Adobe's Distiller and the open source PdfCreator (Figure 6.26, see www.pdfforge.org/pdfcreator for a download).

> **Note**
>
> The PDF specification is very broad and the PdfPlus does not support every feature. If you have trouble during testing of your PDF forms, please contact support via your vendor.

Create and name a new template

Figure 6.27. Template manager before any templates have been created

If no templates have been created on your PdfPlus, the template manager will be empty, with just a link to the template list page (Figure 6.27).

Figure 6.28. The empty template list

To create a new template from the template list page, enter a new template name in the text field and press the add button ( ).

> **Note**
>
> Template names may contain upper and lower case letters, numbers, the dash character and the underscore character.

Figure 6.29. Creating a new template within the template manager

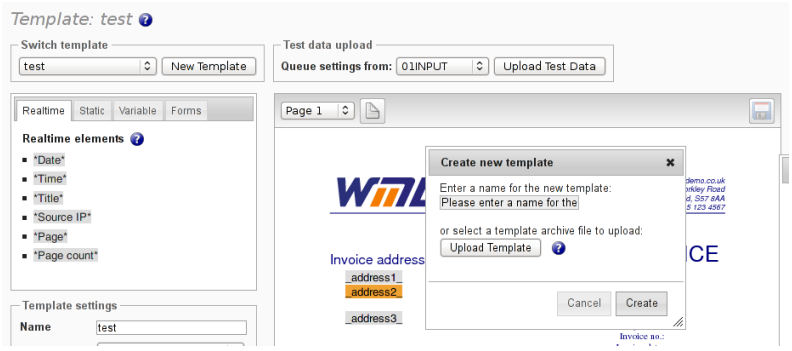

If your PdfPlus already contains at least one template, you can create a new template from within the template manager, by clicking the **New Template** button (Figure 6.29).

Upload and activate PDF forms

Now that you have a template, it is time to upload a form. Select **Data > Template Manager** and choose the **Forms** tab in the overlay elements box (see Figure 6.21).

Figure 6.30. The PDF was uploaded and is now available

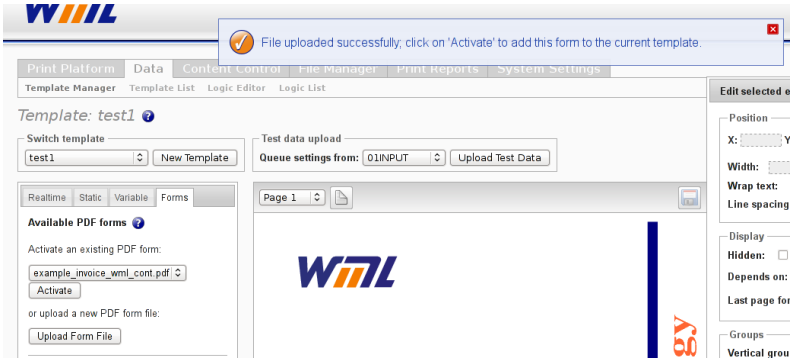

Press the button **Upload Form File** and use your browser's dialog box to select your PDF file. The file will be uploaded to the PdfPlus. The PdfPlus will parse the file to determine if it can successfully overlay the document onto another.

If the PDF form is correctly parsed by the PdfPlus, a success message will read “File uploaded successfully; click on 'Activate' to add this form to the current template” (Figure 6.30).

> **Note**
>
> An uploaded PDF form may be added any number of times to any template on the PdfPlus.

Figure 6.31. The PDF was uploaded but contained an unsupported feature

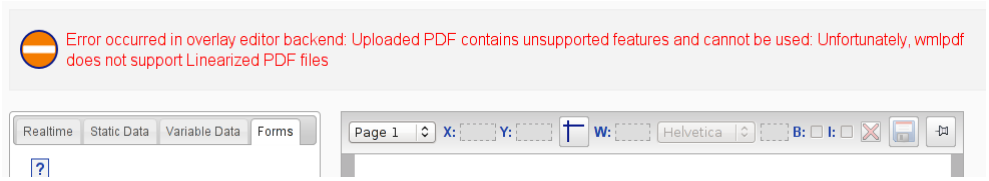

If the PdfPlus encounters a known problem with your PDF form, it will respond with an error message describing the problem (Figure 6.31). In this case, please contact support with details of the problem.

Figure 6.32. Checking the template page

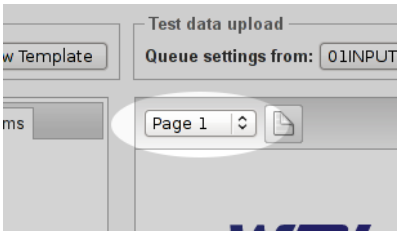

To activate your form, first ensure you are on the page on which the PDF form should be applied. See Figure 6.32 which shows the pages menu. When you activate a form, its first page will appear on the current template page.

Figure 6.33. Activating a form

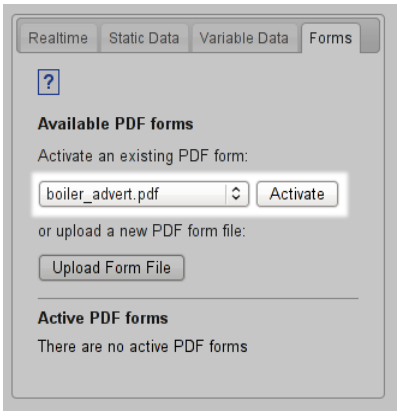

Select the form from the menu highlighted in Figure 6.33. Press the **Activate** button. The form should appear on the canvas, as shown in Figure 6.34.

Figure 6.34. An active form showing on the canvas

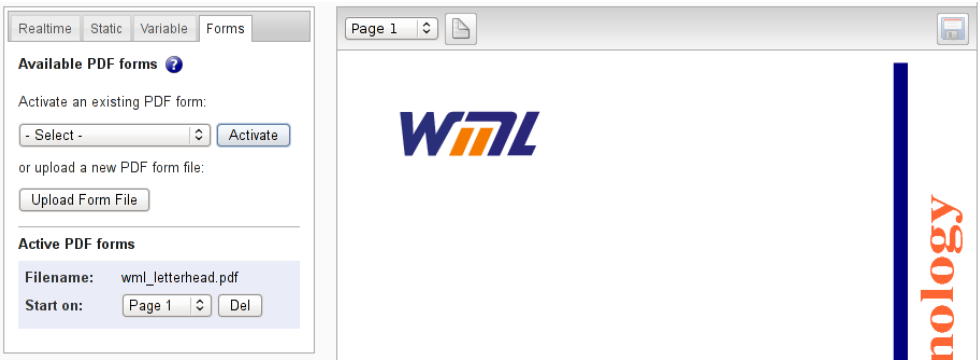

Add text overlays

Now that a form has been added to the template, we'll add a date and time stamp to the example.

Figure 6.35. A date element is added to the template

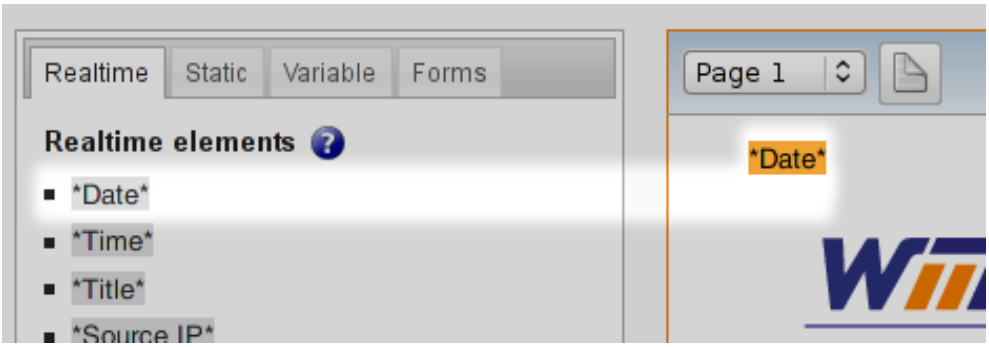

Dates and times are *realtime data* text elements. Select the **Realtime** tab. Drag a date element onto the canvas and position it as shown in Figure 6.35. Do the same for a time text element.

Figure 6.36. Using the snap-to guides to align the text elements

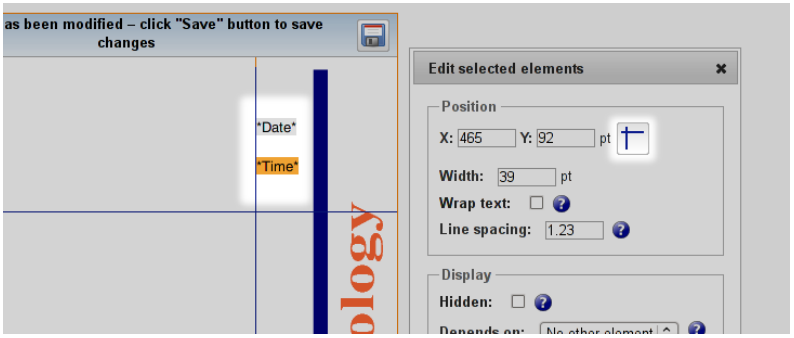

The template manager has a snap-to guide to help you align your text elements. Click the snap-to button ( ) to turn on the guides. Drag the elements towards the guidelines and they will snap into position (Figure 6.36).

Figure 6.37. Moving multiple elements

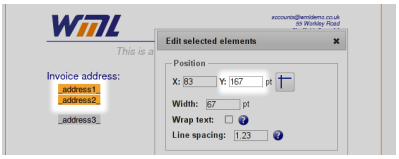

You can select multiple elements and move them together. Click on the first element to select it. Now press the shift key and click on the second element. Both are now selected as shown in Figure 6.37.

To move the selected elements you can edit the text boxes in the element editor; in Figure 6.37 the Y position is being modified. Alternatively, use the cursor keys to "jog" the position of the selected elements.

When you are satified with the position of the text elements, press the save button to write the template to the PdfPlus.

Review template settings

Figure 6.38. The template settings box

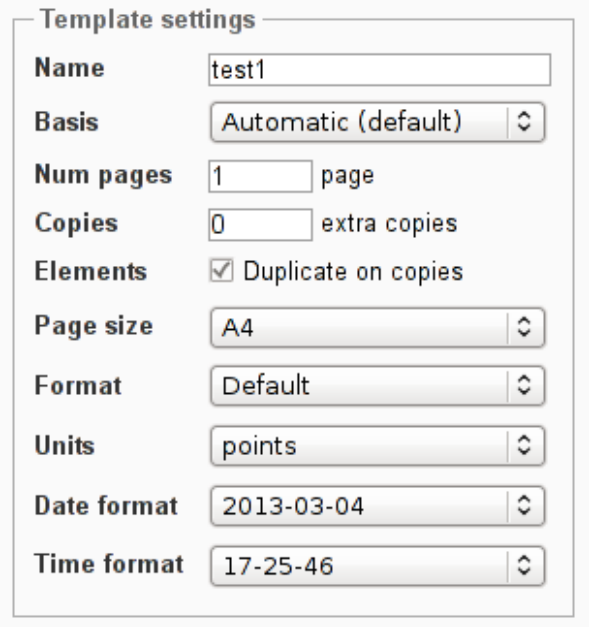

Before we test a print job, refer to the **Template settings** box (Figure 6.38). Check that the **Page size** is set to the same size as your input print data[1]. If the page size here, and the page size of your input print data differ, then text elements may appear out of position when printed.

Finally, choose the format in which you would like the date to be printed, and the format for the time.

Upload and test

To find out how your input print data will be processed by the new template, use the test upload feature in the **Test data upload** box.

[1] Where applicable; if your input print data is structured text data such as CSV or XML formatted data, then choose the page size in which you wish the document to be created.

Select your new queue (here we called it "newQueue") from the menu below **Queue settings from:**. Press the **Upload Test Data** button and use your browser's file dialog to select your input print data.

Figure 6.39. The test print job displaying in a new brower tab

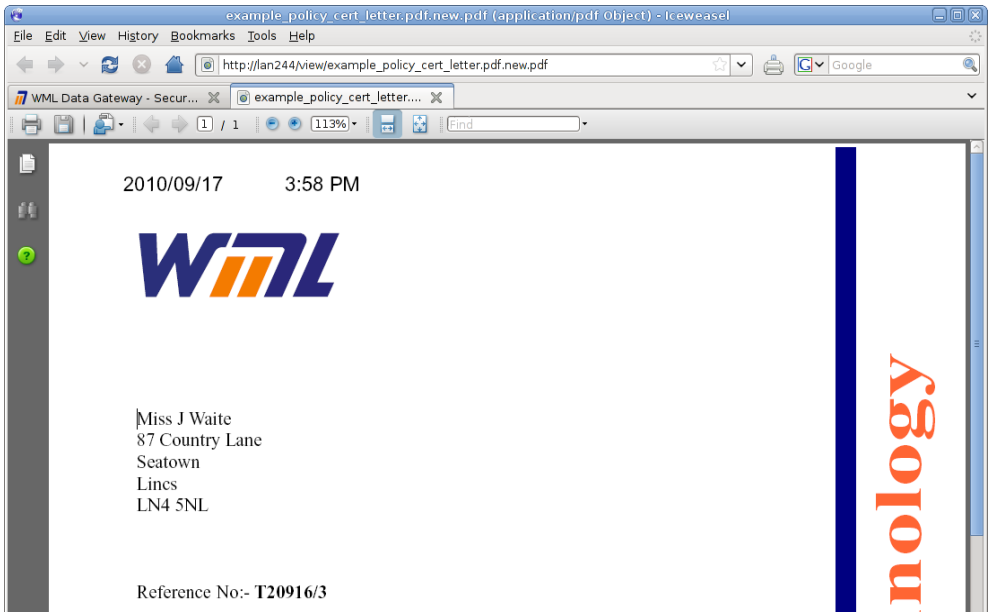

The test print data will be uploaded onto the PdfPlus, it will be processed through the current template, using any queue-specific settings from "newQueue" which will generate an output PDF file. Your browser will be re-directed to open the PDF file in a new tab or window (Figure 6.39).

> **Note**
>
> It's important to allow popups in your web browser. If you upload a test file, and no output PDF document appears this may be the cause.

The example shown in Figure 6.39 shows the input print data (a letter to a Miss J Waite) merged with the PDF form (a WML letterhead) and the date and time realtime data elements, which have been printed as the date and time when the test data example was processed.

Bind the template to the queue

When you are satisfied with the PDF documents generated with your template, you are ready to print a document through your queue. Before printing, you must "bind" your template to the queue.

Enter the **Print Platform > Queue Config** page and switch to your new queue.

Figure 6.40. Binding a template to the queue

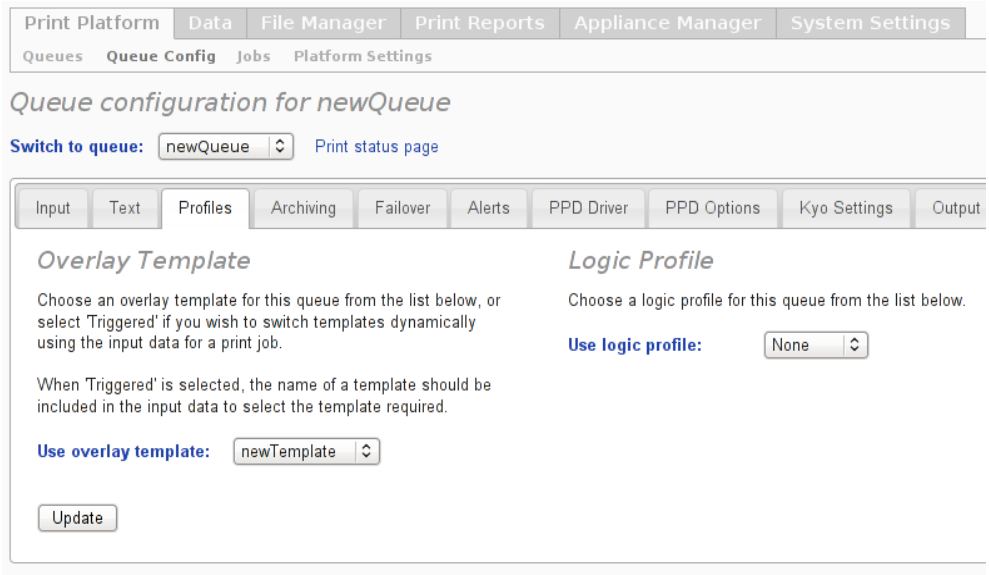

Select the **Templates** tab (Figure 6.40). The menu labelled with **Use overlay template** selects the template which will be applied to print jobs passing through this queue. Select your new template[2] and press the **Update** button.

Send your print data to the PdfPlus using your preferred method (LPD, IPP, raw, hotfolder or HTTP POST) and verify that the document prints as expected.

6.7. Using structured input data

Structured input data is print data sent as a series of specific data items, the value of which may vary between print jobs, but which are to be printed in a pre-defined manner. For example, a database of sales records can be used to produce a set of data relating to each sale. Each data set can then be printed in the form of an invoice, by specifying the position of each data item on a blank invoice template.

The format for structured input data can be XML, comma-separated values (CSV) or Jetforms compatible data. XML data is structured as explicit key-value pairs (and may contain `Job` and `Page` sections), while CSV data may be either a simple ordered list (separated by line breaks, commas or another chosen separator character), in which case the position in the list of an item is the means by which that item is specified, or key-value pairs separated by line breaks.

[2] For greater flexibility you may wish to trigger your template from words in your input print data. See Chapter 7 for more information.

> **Note**
>
> The structured input data code has been designed to make it straightforward to develop additional, customer-specific input data formats.

The type of structured data is detected automatically by the PdfPlus as the print data is received and the data is laid out according to the rules defined in the overlay template assigned to the local queue or corresponding to a trigger value contained in the input data.

XML data is preferred over CSV, as although it appears a little more complex than simple CSV data, the layout is more flexible and errors in the input data are more likely to lead to *no* output rather than *incorrect* output. It is easier to detect that there is an error in the system if no document appears!

XML input data

In XML input data, the first line should be an *XML declaration*, the presence of which tells the PdfPlus that the data is in XML format. The PdfPlus looks for sections and key-value pairs within a parent XML element named `WmlXmlData`.

A single XML input data file may contain the variable data for either a single document, or a batch of many. The input data may be given an overall batch reference, with a `variable` element with `key` attribute equal to "WmlBatchRef" (the value of WmlBatchRef is encoded in the Info dictionary of the generated PDF).

> **Note**
>
> Within the context of structured input data, the term "job" is equivalent to "document" or "document data-set".

Alongside the batch reference, it is also possible to specify the default "WmlTemplate" for all documents in the batch. These two variables are shown in Example 6.2.

Example 6.2. First lines of some XML input data

```
<?xml version="1.0"?>
<WmlXmlData>

  <!-- Batch reference string -->
  <variable key="WmlBatchRef" value="MyInputBatch"/>
  <!-- Batch-default template -->
  <variable key="WmlTemplate" value="myTemplate"/>
```

Job and Page elements

Within the parent `WmlXmlData` element, each document data-set, or "job", should be enclosed within a `Job` element. Each job should be given an `id` attribute, which is the job number, and should start at 1 (rather than 0).

> **Warning**
>
> If the input data contains more than one job, then it is necessary to specify an output queue to process the output documents built from the input batch. This is set in **Print Platform > Queue Config > Text**.

Inside the job element, key-value pairs can be specified either to apply to the entire document, or to each page of the document. Key-value pairs are stored in `variable` elements, each of which defines a single piece of text to be placed on the print job. Each `variable` element should have a `key` attribute and a `value` attribute. The value will be placed and formatted on the document according to the location, font-face and font-size which has been assigned to the key in the Template Manager.

The special key-value "WmlTemplate" can be specified to override any batch-default template for the job. The key-value "WmlDocRef" is used to provide a unique document identifier, which is encoded into the Info dictionary of the generated PDF. "WmlDocRef" is used as the main document identifier within the Document Content Control module (see ???).

If your document needs to display a different value for a given key-value pair on each page, then you should separate the job with `Page` elements. Each `Page` element must contain an `id` attribute, specifying the page number. If you use `Page` elements, you can use the element with id equal to 0 for variables which apply across the job, as shown in Example 6.3.

Example 6.3. Page elements within a job with Page ID 0

```
<Job id="1">
  <Page id="0"> <!-- Page ID 0 means "this isn't a page,
                    it's for the whole job -->
    <variable key="WmlDocRef" value="Report10002"/>
  </Page>
  <Page id="1">
    <variable key="invoice_num" value="125"/>
  </Page>
  <Page id="2">
    <variable key="invoice_num" value="126"/>
  </Page>
</Job>
```

Example 6.4 shows the document wide key-value directly within the `Job` element. Example 6.4 is exactly equivalent to Example 6.3. In both cases, this information is to be used to build a report on which the value for "invoice_num" is "125" on page 1 and "126" on page 2. The unique document reference for the job is "Report10002".

Example 6.4. Page elements within a job

```
<Job id="1">
  <variable key="WmlDocRef" value="Report10002"/>
  <Page id="1">
    <variable key="invoice_num" value="125"/>
  </Page>
  <Page id="2">
    <variable key="invoice_num" value="126"/>
  </Page>
</Job>
```

If the template is specified with "WmlTemplate" for the entire document, then the template is expected to have as many pages as there are `Page` elements in the XML, or the template should have an **Others** page.

> **Note**
>
> See Section 6.5 for more detail about the Template Manager. The **Others** page in a WML Template specifies the layout for all pages which are not uniquely laid out. For example, you may specify a first page layout for a letter, a second page layout for some terms of business, then an **Others** page for all continuation pages after page 2.

You can use templates as "page templates", rather than as document-wide templates. In this mode, you specify a WmlTemplate for each page in the job. Only the first page of each template is applied to the page of the document.

Example 6.5 shows a full, annotated example of an XML input batch file for multiple jobs.

Example 6.5. Full XML example

```
<?xml version="1.0"?>
<WmlXmlData>

  <!-- Batch reference string -->
  <variable key="WmlBatchRef" value="Input Batch Name"/>
  <!-- Batch-default template -->
  <variable key="WmlTemplate" value="someTemplate"/>

  <Job id="1">
    <Page id="0"> <!-- Page ID 0 means "this isn't a page,
                       it's for the whole job -->
      <variable key="WmlDocRef" value="Doc 1"/>
    </Page>
    <Page id="1">
      <variable key="key1" value="something1"/>
    </Page>
    <Page id="2">
      <variable key="key1" value="something1.1"/>
    </Page>
  </Job>

  <Job id="2"> <!-- NB: No page elements in this one,
                    page is implicitly equal to 0. -->
    <variable key="WmlDocRef" value="Doc 2"/>
    <variable key="WmlTemplate" value="overrideTemplate"/>
    <variable key="key1" value="something2"/>
  </Job>

  <Job id="3"> <!-- This job shows the use of per-page templates -->
    <variable key="WmlDocRef" value="Doc 3"/>
    <Page id="1">
      <variable key="WmlTemplate" value="CoveringLetter"/>
      <variable key="name" value="Bob"/>
    </Page>
    <Page id="2">
      <variable key="WmlTemplate" value="TermsOfBusiness"/>
      <variable key="date" value="2011-01-03"/>
    </Page>
  </Job>

</WmlXmlData>
```

Previous XML data format

An example of the older format for XML structured input data is shown in Example 6.6. This format lacks the flexibility to specify job and page elements, but the PdfPlus will still accept this data.

Example 6.6. XML input data - old format

```
<?xml version="1.0" ?>
<!DOCTYPE WmlXmlData >
<WmlXmlData>
  <variable key="WmlTemplate" value="sales_contract"/>
  <variable key="WmlDocRef" value="REF123456"/>
  <variable key="copyLabel" value="Customer Copy"/>
  <variable key="refNo" value="7101-E9418"/>
  <variable key="company" value="CFG SUPPLIES"/>
  <variable key="address1" value="1 MINNS WAY"/>
  <variable key="address2" value="FAIRFORD RETAIL PARK"/>
  <variable key="address3" value="WESTFIELD, YORKS W1 9GT"/>
  <variable key="telephone" value="Tel: 01234 567890"/>
</WmlXmlData>
```

Character encoding and special characters

The PdfPlus supports the following character encodings for XML input data:

- UTF-8
- UTF-16 (both little and big endian)
- ISO-Latin-1 (ISO-8859-1)
- ASCII

The default encoding is UTF-8. To use one of the other supported encodings, specify the encoding in the XML declaration as shown below.

```
<?xml version="1.0" encoding="iso-latin-1" ?>
```

A number of characters have special meaning in XML data. If you wish to include any of these characters in the print data you must use the corresponding character sequence shown in the following table.

Table 6.1. Special characters in XML

<	<
>	>
'	'
"	"
&	&

> **Note**
>
> The characters used for data keys should be confined to the ASCII character set and should not include any of the special characters above.

CSV data

Comma-separated format is simpler than XML, but a little more restrictive and also prone to error. The main drawback to using comma-separated values is that the character used as separator cannot be used within the values.

This restriction aside, separating values using commas, spaces or new-line characters is often adequate, and many applications generate output data in this format.

> **Note**
>
> To ensure that the PdfPlus consumes CSV data as structured data, rather than plain text, set **Treat input text** to "as structured data" in **Print Platform > Queue Config > Text**.

PdfPlus accepts data in two forms of CSV; "ordered" and "key-value". Ordered data is simply a list of values, separated by a special character, with no keys in the listing (as described below, the keys are implied). Key-value data is a table of key-value pairs.

> **Note**
>
> The field separator character can be specified in **Print Platform > Queue Config > Text**.

PdfPlus will auto-detect whether the data has been arranged as ordered CSV or in key-value pairs. Care must be taken to ensure that the separator character does not appear in either the keys or the values, otherwise unexpected and confusing results will be seen.

Ordered CSV

Ordered CSV is a simple list of values, one value following the next, separated either by new-lines or by the separator character. The new-line can be either the line-feed (0xa, LF or \n) character used in Unix (including Linux and Mac) systems or the DOS carriage-return/line-feed pair (0xd 0xa, CRLF or \r\n).

The ordered list has implied keys. Internally, PdfPlus assigns this list the keys 0, 1, 2, 3, and so on. The ordered list can be separated by new-lines, with a new value on each new line, as shown in Example 6.7 or by a user-specified character such as a comma (default) or semi-colon (see Example 6.8). In both examples, the implied key "0" is "sales_contract", key "1" is "Customer Copy" and so on.

Example 6.7. Column list CSV data

```
sales_contract
Customer Copy
CFG SUPPLIES
1 MINNS WAY
FAIRFORD RETAIL PARK
WESTFIELD, YORKS W1 9GT
Tel: 01234 567890
7101-E9418
```

Example 6.8. Row list CSV data

```
sales_contract;Customer Copy;CFG SUPPLIES;1 MINNS WAY;FAIRFORD RETAIL PARK;
WESTFIELD, YORKS WS1 9GT;Tel: 01234 567890;7101-E9418
```

Key-value CSV

The format of key-value CSV data is shown in Example 6.9. The data is laid out as a table. The columns of the table are delimited by the user-specified separator character and the rows of the table by new-lines. The first, left-hand column contains the keys, the second column contains the values.

In the example, the key “address1” has the value “1 MINNS WAY” and key “telephone” has the value “Tel: 01234 567890”.

> **Note**
>
> Unfortunately, it is *not* possible to arrange the key-value table as two rows instead of two columns.

Example 6.9. CSV key-value data

```
WmlTemplate,sales_contract
copyLabel,Customer Copy
company,CFG SUPPLIES
address1,1 MINNS WAY
address2,FAIRFORD RETAIL PARK
address3,WESTFIELD, YORKS W1 9GT
telephone,Tel: 01234 567890
refNo,7101-E9418
```

Specifying multiple jobs

To specify multiple jobs in a single CSV file, a string can be specified as a job separator. Any lines of the CSV input containing this string will cause one job to end and the next to begin. The job separator is specified in **Print Platform > Queue Config > Text**. If the input data contains more than one job, then it is necessary to specify an output queue to process the output documents built from the input batch. This is also set in **Print Platform > Queue Config > Text**.

Character encoding and special characters

Data should be encoded as ASCII or ISO-Latin-1 (ISO-8859-1). For key-value data, the key should be confined to the ASCII character set and should not include any of the characters <, >, ', ", or &.

Jetforms format input data

The PdfPlus will accept incoming data in Jetforms format. For more information, please contact WML.

Template special keys

For all data formats there is a key (or position) that is reserved for specifying the template to be used for the current print job. In the following example, the queue has been set up to select the template by the "Triggered" option (the section called "Overlay templates and logic profiles"). If the value specified by the key corresponds to the name of an existing template then that is what will be applied to the current print job. In this way, structured print data passed through a single queue can be printed in a variety of ways according to which template is selected.

For key-value XML or CSV input data, the reserved key used to specify the template is "WmlTemplate". If your input data is ordered CSV, then the first entry (implied key "0") is used to specify the template. The example listings above all show the use of the reserved trigger key to specify the overlay set "sales_contract"

Also reserved are the keys "WmlDocRef" and "WmlSortKey". If either is present, then its value will be inserted into the info dictionary of the PDF file as a metadata item. It is intended that the customer or document reference number or identifier would be placed into "WmlDocRef".

The Document Content Control module makes use of the WmlDocRef for automatic document collation. WmlSortKey is used for any information which should be used for pre-output document sorting. For example, this key could hold the destination postcode for the document, and that information would then be used to mailsort the documents prior to printing.

> **Note**
>
> Unlike the WmlTemplate key, the WmlSortKey and WmlDocRef keys can *only* be specified in key-value data (XML or key-value CSV).

6.8. Using data groups

This section provides examples of using the data groups feature of the Template Manager.

Letter with variable content

Consider creating a template for a letter, the contents of which are not known in advance. In particular, the main body of the letter may comprise several paragraphs, the number and contents of which will affect the length of the resulting document.

We can use a *data group* element to indicate an area of the canvas to be filled with the paragraphs in the main body of the letter. A set of variable data elements will be specified for the group, and as many of these as occur in the input data will be added to the data group area, and laid out according to the data group's parameters.

Figure 6.41. Example showing a data group for the body of a letter

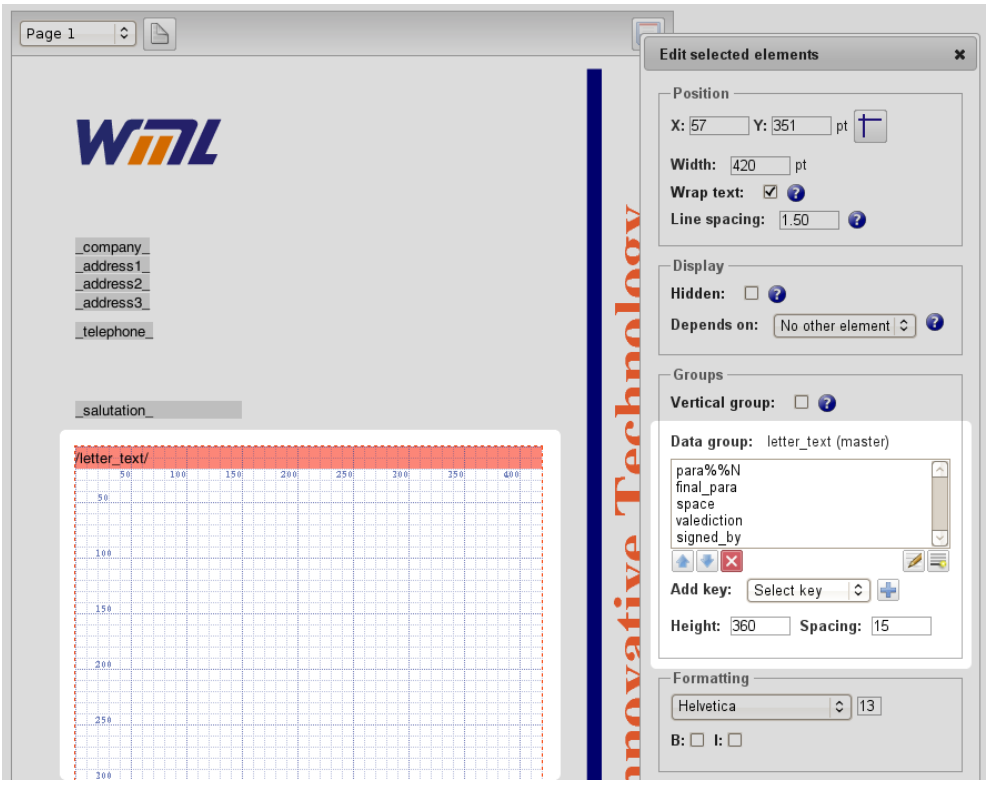

Figure 6.41 shows the canvas for our letter template, with a data group called “letter_text” to hold the main body of the letter placed beneath the “salutation” field.

The data groups controls in the element editor are highlighted in Figure 6.41. These fields are displayed only when a data group element has been selected. The name of the data group is given, followed by the word “master” in parentheses. This indicates that this is the master data group, for which we'll specify the variable data elements to be displayed.

The multiple select box has been filled with several keys. These indicate the variable data to include in the group. Entries may be added to the list using the “Add key” dropdown below, which allows you to select from a list of example variable data keys or to specify the key text manually. The controls below the list can be used to delete or re-order the entries.

The first entry in the list is **para%%N**. The special tag **%%N** is used to indicate sequentially numbered elements. In this example, the entry indicates variable data elements with keys para1, para2 and so on. When an input job is processed, variable data elements with the corresponding keys will be added to the data group until the key for the next “N” is not found. At this point the next element in the list, if there is one, will be added.

The remainder of the keys displayed in this example are normal variable data keys. These are used to append the valediction and so on to the body of the letter.

Below the key list are fields for the height and spacing of the data group. The “Spacing” field governs the spacing between individual variable data elements in the group. Data group formatting also depends on width, text wrapping and line-spacing settings, which are found in the “Position” section of the element editor. Text wrapping is enabled by default for data groups.

Figure 6.42 shows the result of processing an input data file containing two variable data entries for the paragraphs in the main body of the letter. The two short paragraphs are added in sequence, followed by the closing lines of the letter, as specified in the variable data key list.

Figure 6.42. Output from the letter template

Now we need to handle letters for which the amount of text in the main body of the letter overflows the area specified for the master data group element, as in a longer letter with multiple paragraphs in the body text.

When a master data group is added to a template page, a new entry appears in the “Pages” dropdown menu (Figure 6.43). This entry is the *continuation page* for the data group, a tem-

plate page that can be set up to display data that overflows from the data group area on the normal template page. For page 1 of the template, the new entry will be “Page 1 cont”.

Figure 6.43. Pages menu showing continuation page entry

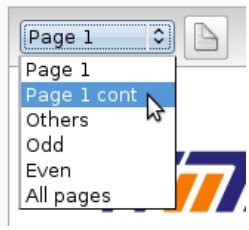

The contination template page allows you to set up a page with a different layout to the original page on which to continue the display of elements in the data group.

> **Note**
>
> If no overlay elements are added to the continuation template page, a copy of the original page will be used as a continuation page instead.

To set up the continuation page to display the overflow from a data group, drag a copy of the data group from the “Existing data groups” list onto the continuation tempate page.

This creates a *continuation data group* for the original master data group. The variable data key list cannot be edited for this continuation group; the elements to be displayed are as specified for the master data group. However, the position, size and spacing for the continuation group may be changed as required.

You can add any other overlay elements, such as forms, static text or realtime data, to the continuation template page.

Figure 6.44. Letter template pages - original page (back) and continuation page (front)

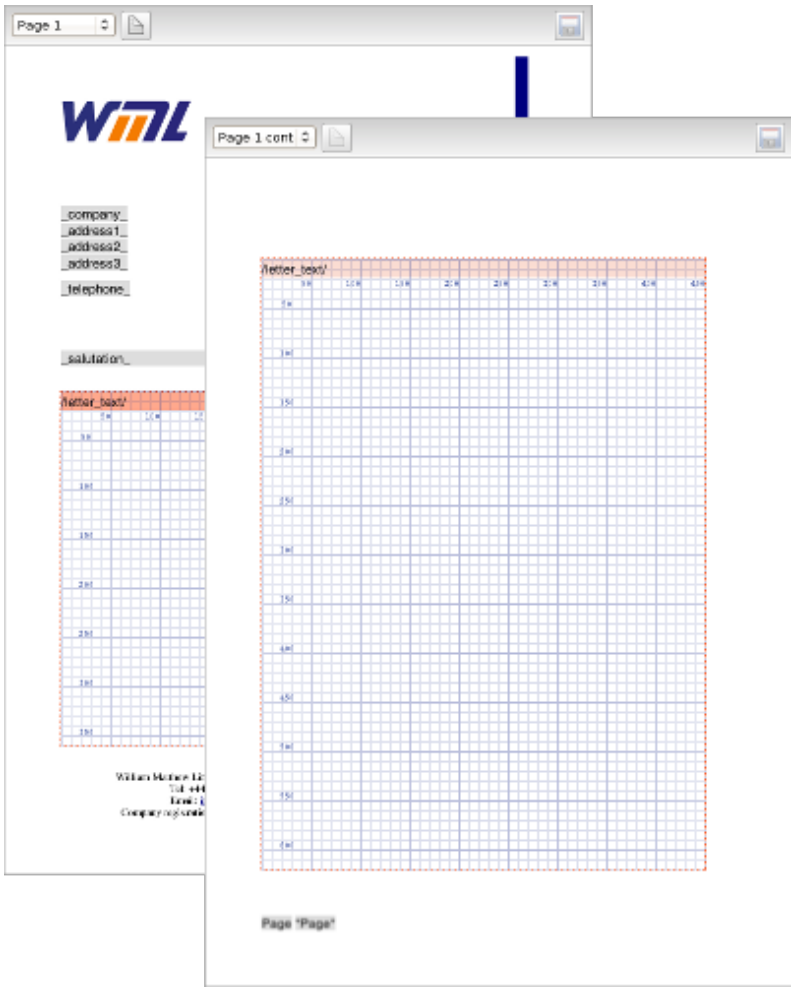

Figure 6.44 shows the continuation page for our example set up to display the remainder of the letter text on a plain page, starting at the top margin of the page and continuing down the full length of the page. Two text overlay elements have been added at the bottom of the page to display the current page number. Figure 6.45 shows the result of processing an input file for which the main text of the letter overflows onto the continuation page using this template.

In this case, the input data contains a long paragraph which overflows the space available in the master data group. The paragraph has been stopped where it reaches the bottom of the master data group. The remaining text in the paragraph has been placed on the continuation page, followed by the subsequent variable data elements. The page number text has been inserted with the correct page number.

Figure 6.45. Output from the letter template with overflow to continuation page

Multiple page invoice

In this example we look at using data groups to build a multiple page invoice. We introduce the concept of data group *columns* and show how these can be used to display table-based data.

Consider the invoice template shown in Figure 6.46. At the top of the invoice are standard variable data elements for the invoice address and reference numbers. Below is a table with columns for itemizing the invoice contents. Each item of the invoice will be represented

by several variable data elements, which need to be displayed together as rows in the body of the table.

Figure 6.46. Example invoice template

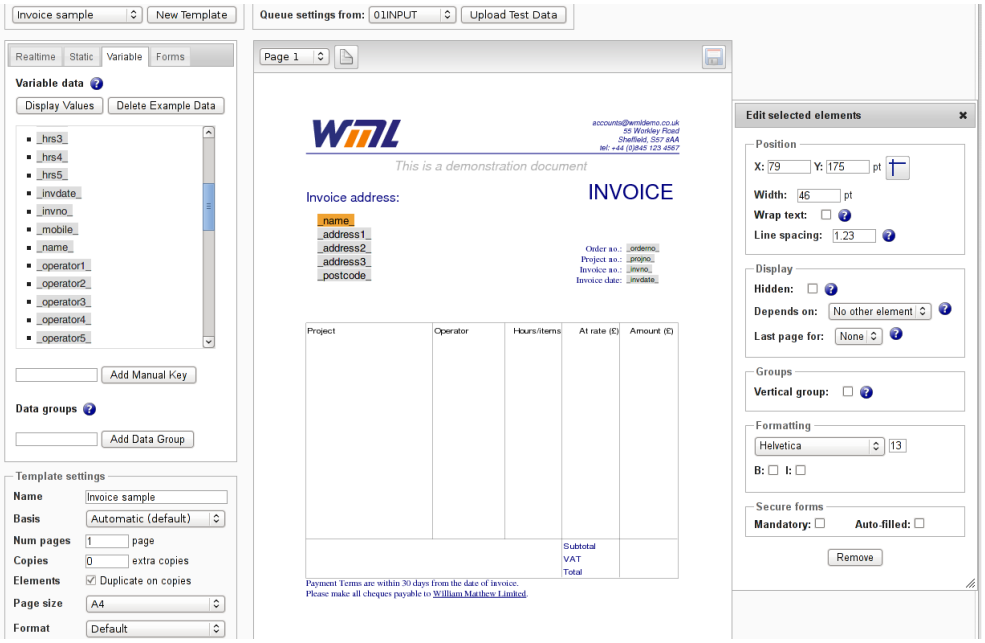

The first step is to add a data group for the invoice items (Figure 6.47). The newly added group has a default width and height, which we'll need to modify so that the data group corresponds to the area of the invoice table. This can be done using the width and height fields in the element editor, highlighted in Figure 6.47. Note that the data group area is displayed with a background showing gridlines. These will help us later with setting up the layout for the data group.

Figure 6.47. Adding a data group to the invoice template

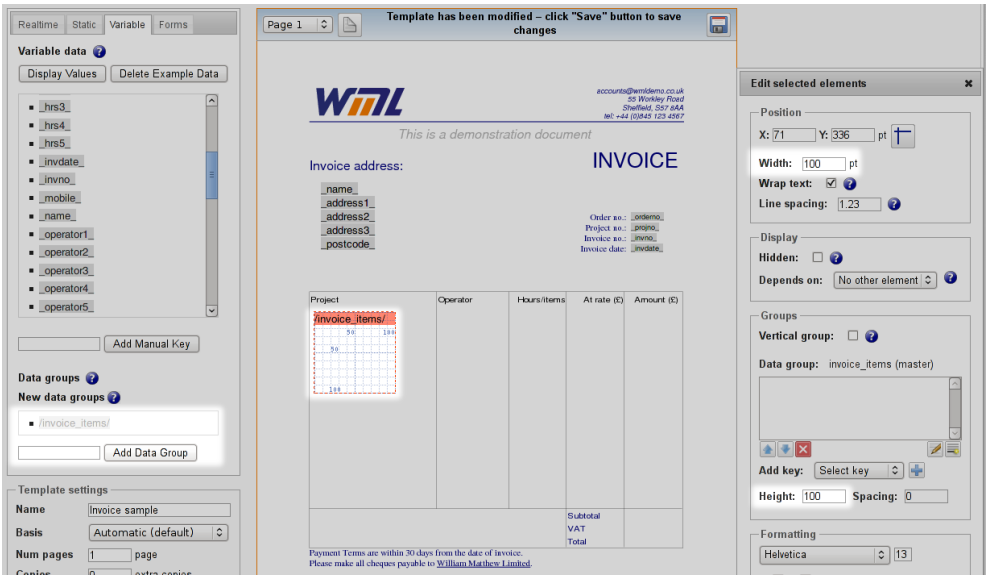

For this example we need to create a row of variable data elements for each invoice item. We'll do this using the “Create data group row” controls, which are accessed using the button at the bottom right of the data group key list (Figure 6.48).

Figure 6.48. Creating a data group row for the “invoice_items” data group

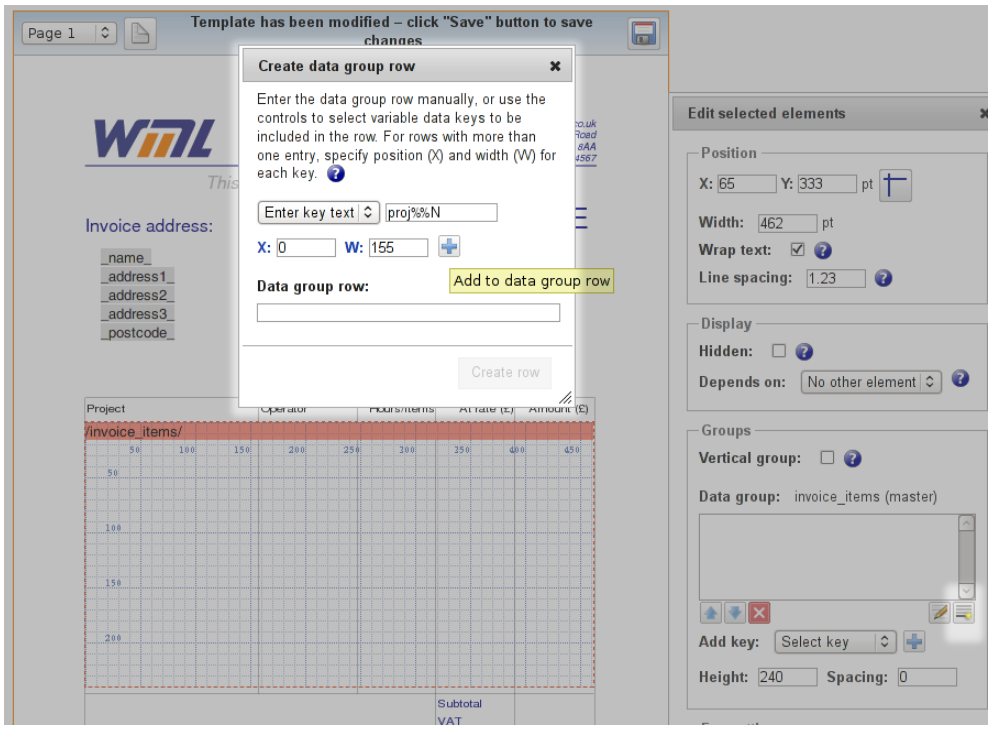

Each data group row may have a number of entries or *columns*. In the simplest case, the row is a single entry, and the only thing that needs to be specified is the variable data key that indicates the data value to insert. For tabular data, the row can be specified as a number of columns, each of which has a position and width within the row.

Figure 6.48 shows the controls for creating a data group row with columns. You can use the variable data key dropdown and the “X” (position) and “W” (width) fields to specify each column, using the add button ( ) to add columns to the row. Alternatively, you can edit the data group row directly in the field below. The format for a data group row with columns is as follows:

```
key1,position1,width1,...,keyN,positionN,widthN
```

The keys may be specified explicitly or using the special tag **%%N** to indicate sequentially numbered elements (see the “Letter with variable content” example). The position is the horizontal distance in PDF units (points) of the column from the left hand edge of the data group. Similarly, the width specifies the width of the column in PDF units. You can use the gridlines displayed for the data group area to calculate appropriate values. When you have set up the required columns, click on “Create row” to add the row to the data group key list.

Figure 6.49 shows a document produced from an invoice template in which a data group row defines entries for the project, operator, hours, rate and amount columns of the invoice table. The entries are lined up in columns as specified in the data group row. Note that the project descriptions are wider than the column to which they belong, so they have been wrapped to fit.

Figure 6.49. Sample output from the invoice template

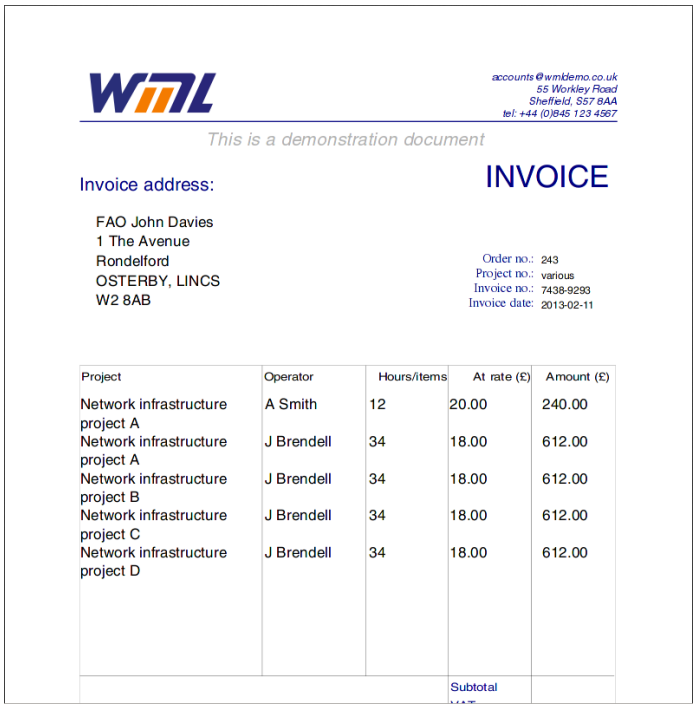

The positions of the columns in Figure 6.49 could be tidied up. To edit an existing data group row, highlight the row in the data group key list and click on the edit button (). This brings up a dialog where you can manually edit the data group row (Figure 6.50).

Figure 6.50. Editing a data group row

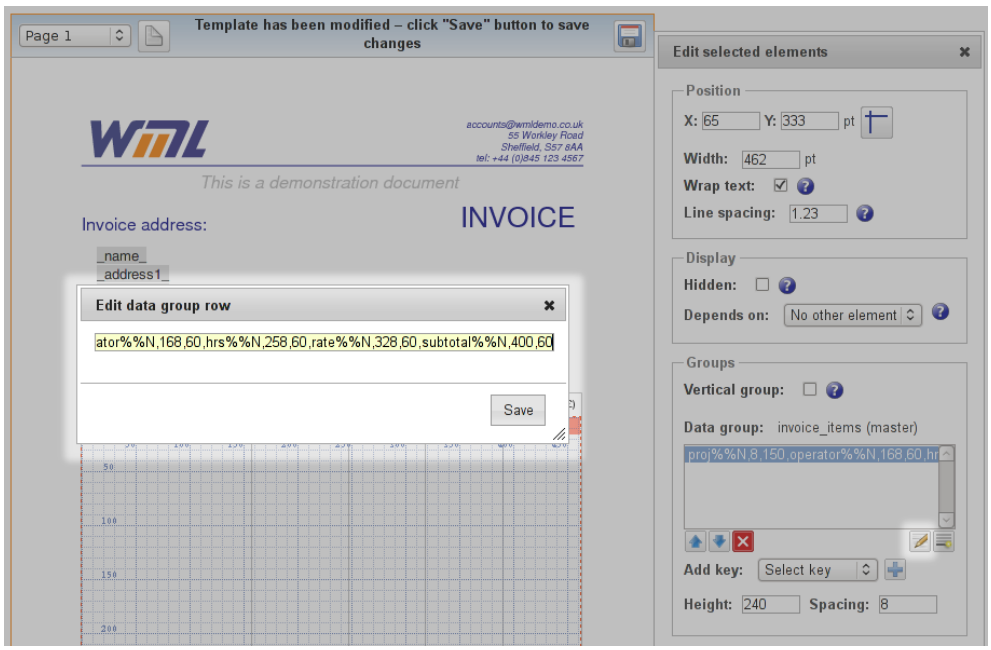

After modifying the row in accordance with the data group row format indicated above, click on “Save” to update the data group key list. Note that you'll also need to click on the save button ( ) to save the changes to the template.

Figure 6.51 shows the same document as before, with the column positions adjusted. The data group spacing, which governs the spacing between the rows, has also been increased.

Figure 6.51. Sample output with columns and spacing adjusted

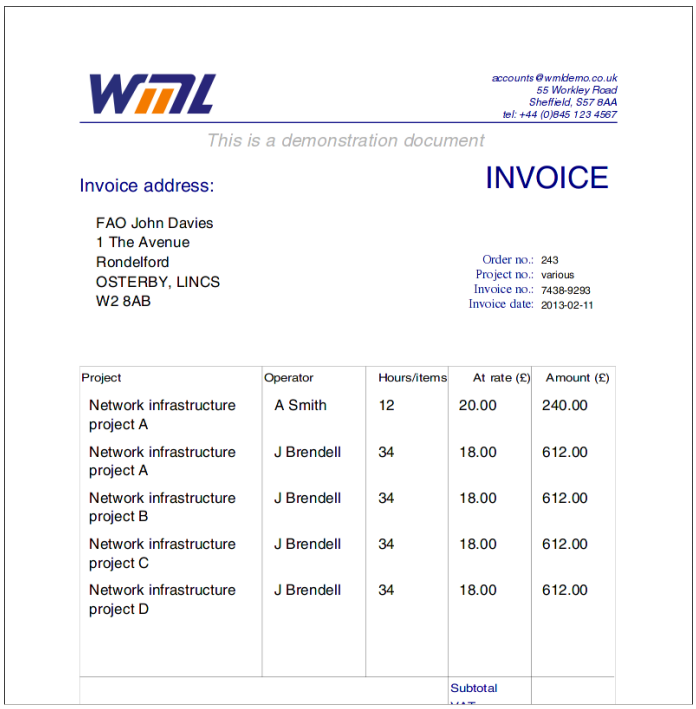

As in the “Letter with variable content” example, the document may need to be continued on one or more extra pages. Figure 6.52 shows a continuation template page for the invoice. This uses a background form with an extended invoice table. A continuation group for the “invoice_items” data group has been added to display any items which overflow from the master data group. The continuation group will use the same data group rows as specified for the master data group, so it will have the same column layout.

Additional variable data fields have been added to display subtotal, VAT and total figures for the invoice. Since these entries should come at the end of the invoice, they have been set as “Last page for” the “invoice_items” data group in the element editor, meaning that they will only be output when the last continuation page has been reached. (If there are no continuation pages, the elements will be output on the original page.)

Figure 6.52. Continuation template page for the invoice

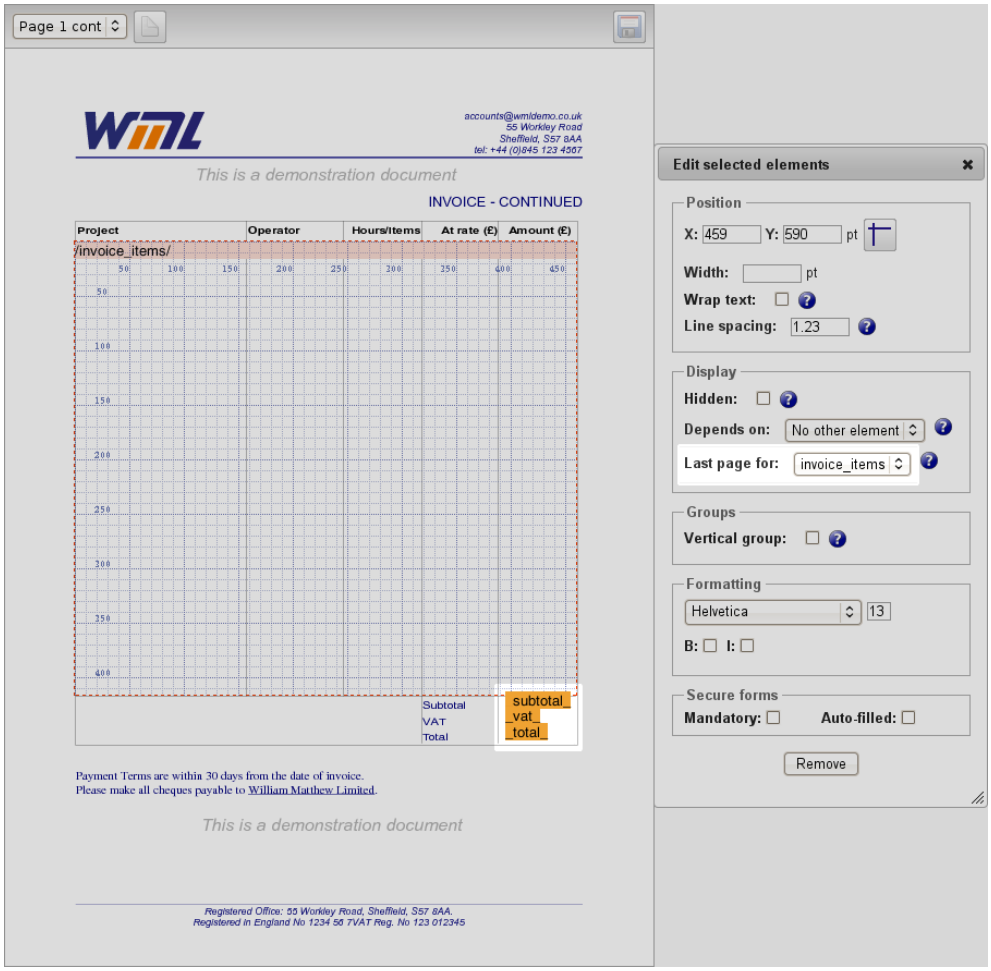

Figure 6.53 shows the result of processing a data file for a longer invoice, which runs on to a second page.

Figure 6.53. Invoice with overflow to continuation page

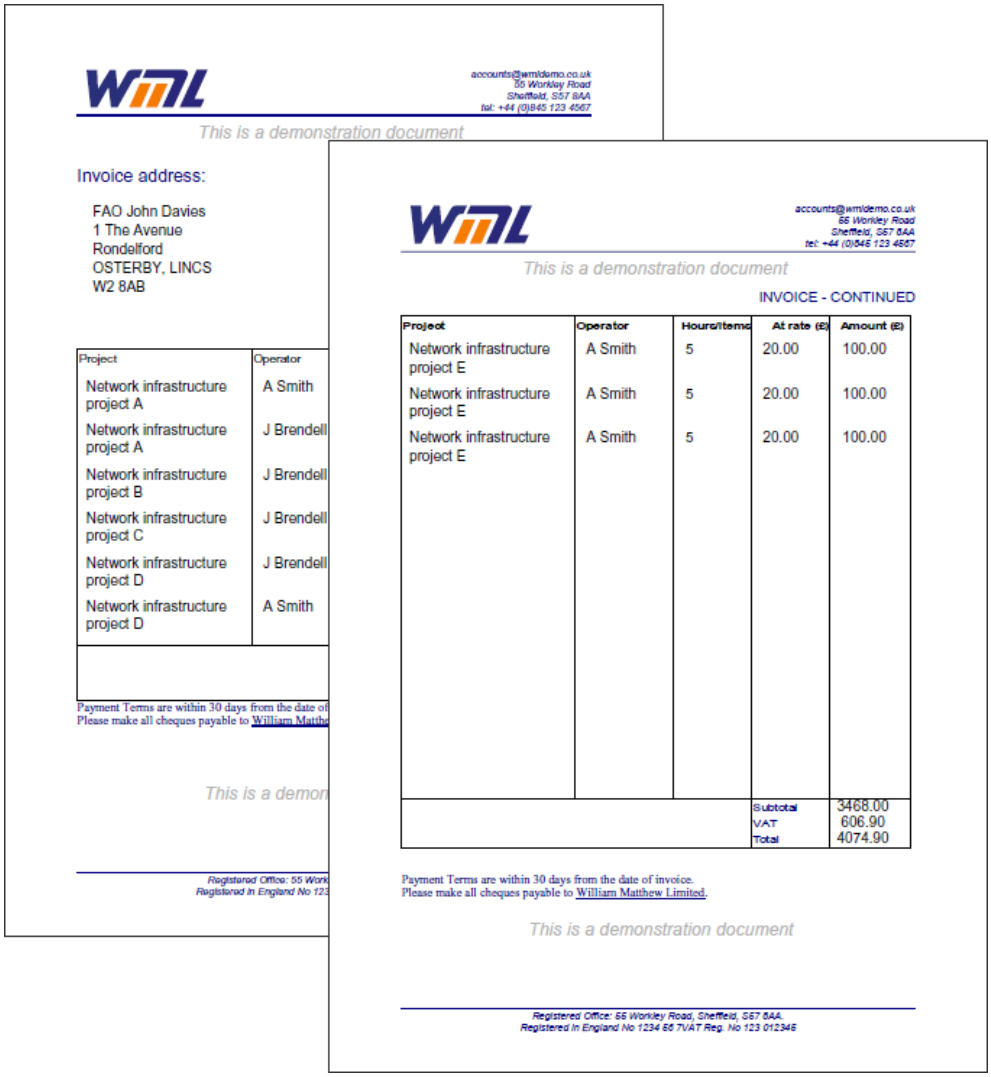

Chapter 7. Logic Profiles

Logic profiles provide a way to trigger data manipulations on incoming print data *based on the content of the data*.

Each logic profile contains a number of *logic rules*. A logic rule contains a number of *logic conditions*. If these conditions are met by the incoming document, then the *logic action* is carried out.

Logic actions include:

- Text search and replace
- Insert a barcode of the searched text at a specified location on the page.
- Trigger a template for the entire document.
- Trigger a PDF form for a page which meets the logic rule's conditions.

Figure 7.1. An empty logic profiles page

Logic profiles, like overlay templates, are managed using a logic profile editor and a logic profile list. To create the first logic profile, go to **Data > Logic List** (see Figure 7.1).

Figure 7.2. A new logic profile

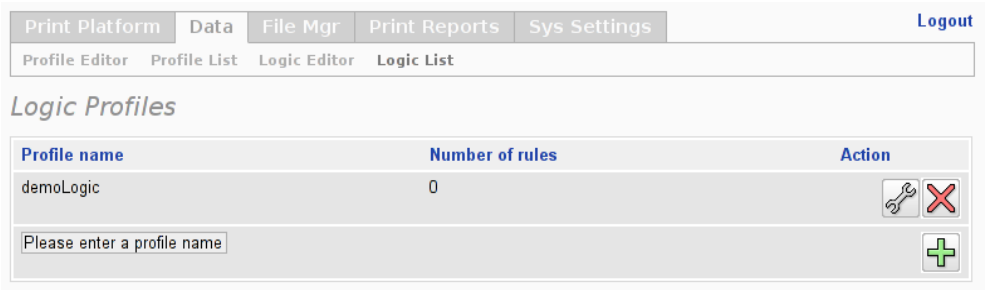

Add a new profile by filling in the logic profile name, and pressing the add button ( ). The new profile will appear in the list, as shown in Figure 7.2. Click on the configure button ( ) to access the logic editor page for the new logic profile. The empty logic profile is shown in Figure 7.3.

Figure 7.3. A new logic profile in the logic editor

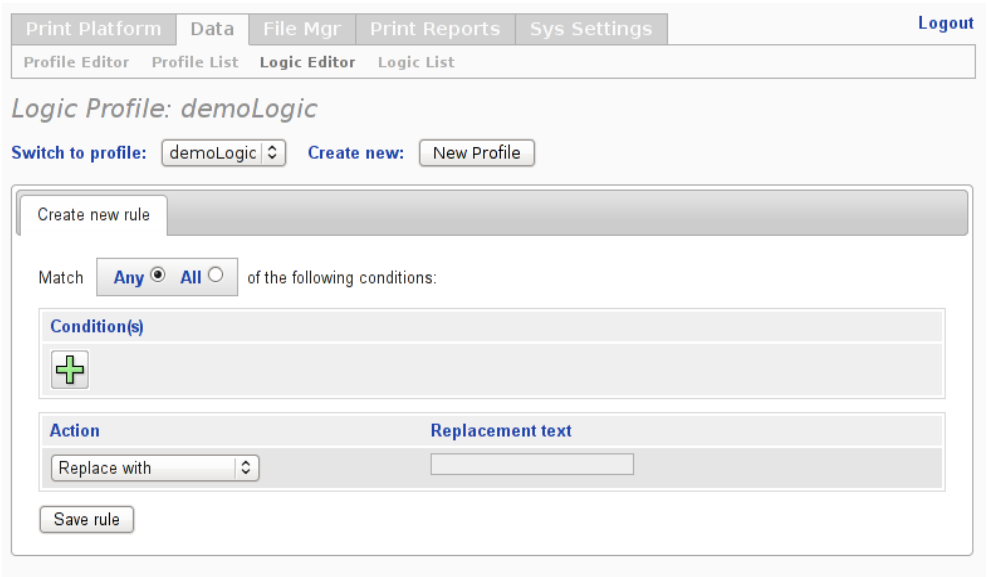

In Figure 7.3, you can see a **Switch to profile** menu and a **New Profile** button, similar to those in the overlay template editor. These function in the same way.

The primary region of the logic editor is the rules box. In Figure 7.3 this shows a single tab headed **Create new rule**. Here you can add conditions to the rule and choose an action.

Figure 7.4. Adding a new condition to the rule

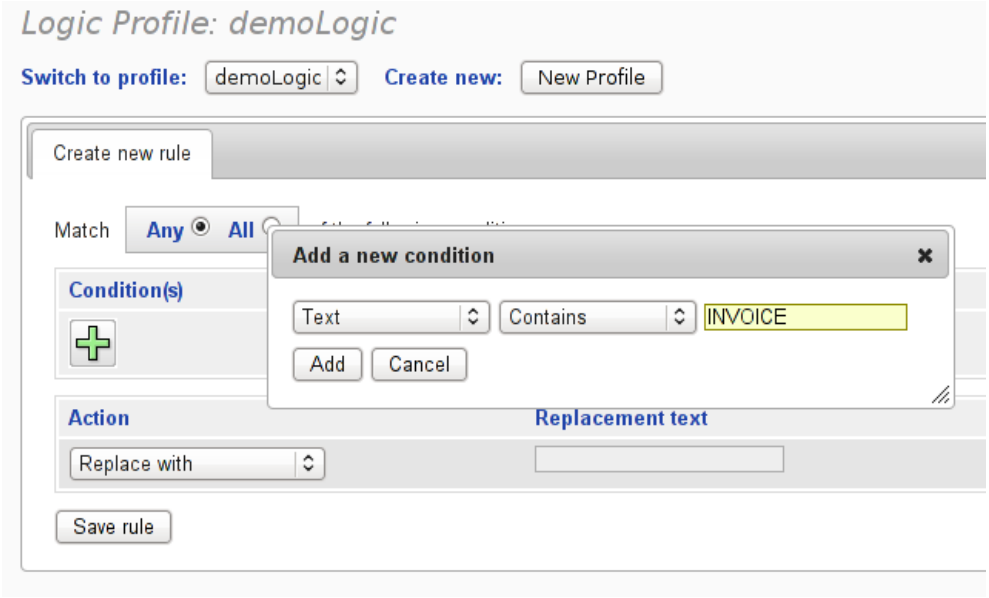

To create the new rule, you must add at least one condition and select the action.

To create a condition, press the add button ( ). A dialog will pop up to create the condition. This is shown in Figure 7.4.

> **Note**
>
> The left menu in Figure 7.4, whose value is "Text" currently has no other values as options. This menu is in place for future use.

Decide whether this condition should be "text contains" or "text does not contain", and enter a match word or phrase. In the example shown, a condition that text should contain the word "INVOICE" is shown.

> **Note**
>
> All text conditions are case sensitive. If you need to match both upper and lower case, use two conditions and specify that **Any** condition must match.

Figure 7.5. Choose an action for the rule

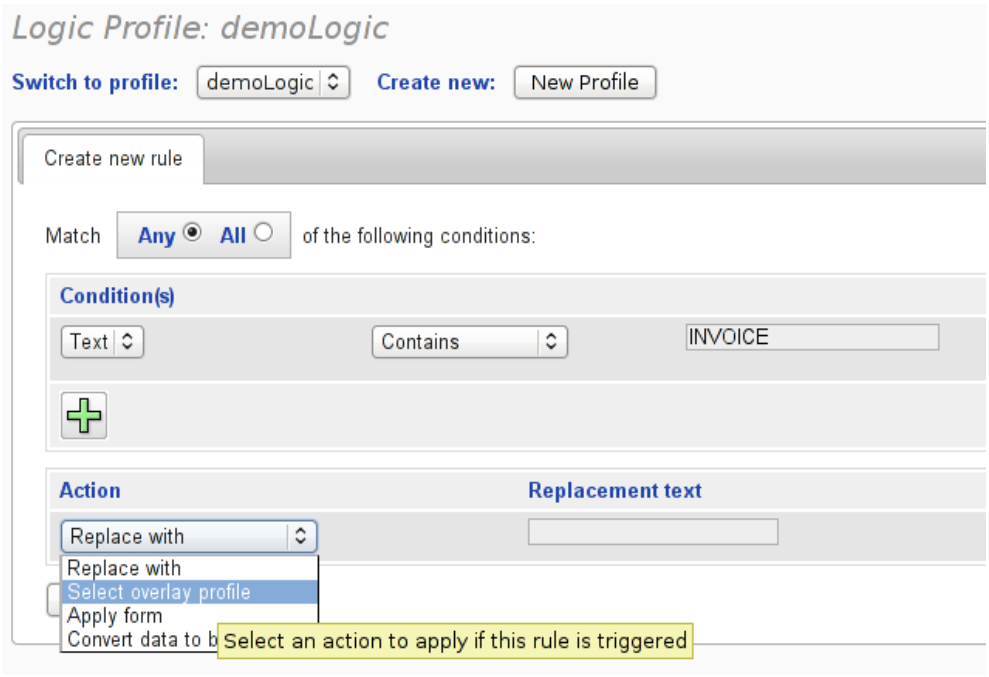

Press the **Add** to add the condition to the rule. You can now add additional conditions and use the **Any/All** radio buttons to specify whether all of the conditions must match or whether it is sufficient for any condition to match.

Choose an action from the **Action** menu, as shown in Figure 7.5. The parameter for the action will depend on the action selected (Figure 7.6 and Figure 7.7).

Figure 7.6. Apply form - selecting a form

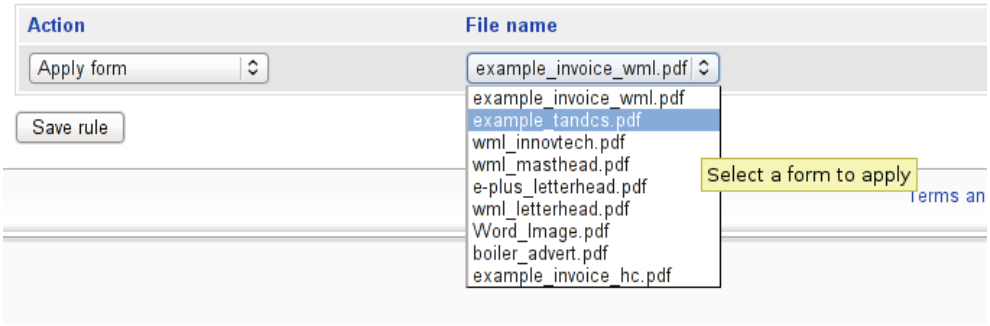

Figure 7.7. Select overlay template - available templates

Figure 7.8. The newly created logic rule

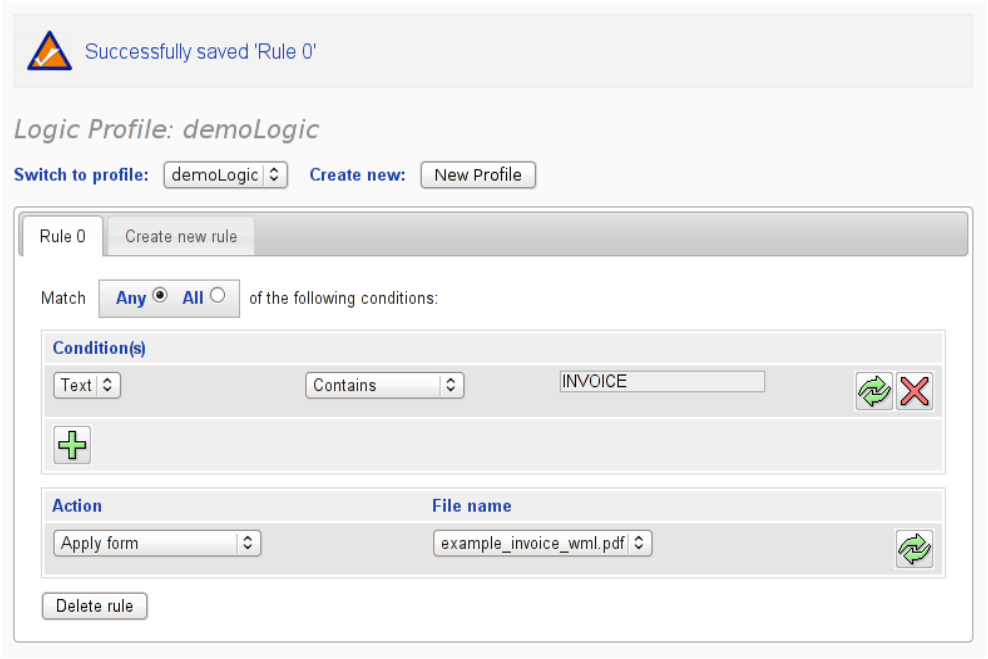

When you are satisfied, press the **Save rule** button. The tab will be re-named to **Rule0**, with a **Delete rule** button and a new **Create new rule** tab will appear. This is shown in Figure 7.8

The example rule will apply a PDF form called “example_invoice_wml.pdf” to any page of the input print data which contains the text “INVOICE”.

To make use of the logic profile, it must be *bound* to a queue in the same way that a template can be associated with a queue.

Figure 7.9. Binding a logic profile to a queue

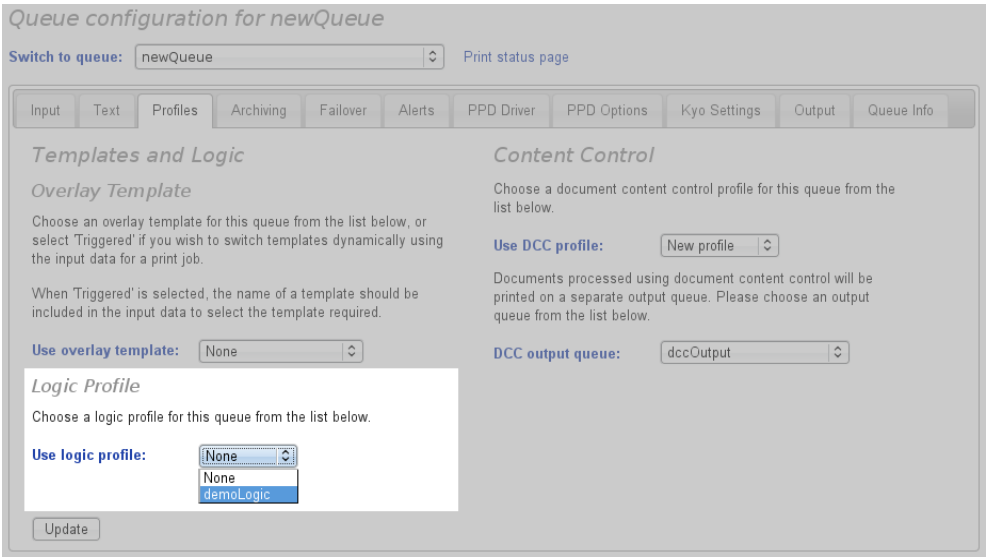

Access the queue configuration page at **Print Platform > Queue Config**. Select your profile from the the menu (Figure 7.9) and press **Update**. Print jobs on this queue ("newQueue") will now be processed through the logic filter.

Chapter 8. File Manager

To access the built in file manager, select **File Manager** from the top level menu. Your browser will display a page similar to Figure 8.1.

Figure 8.1. The PdfPlus's file manager

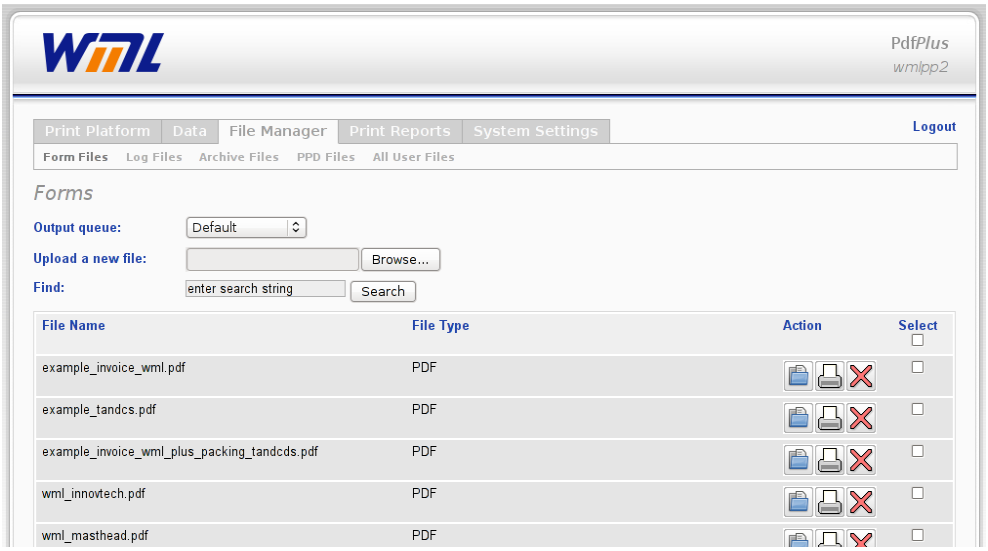

The File Manager has a number of sections listing files relating to different features of the PdfPlus.

- Form Files: User-uploaded PDF files for use as forms or digital stock.
- Log Files: Daily system logs stored on a network share or on internal storage.
- Archive Files: Archived print jobs, in PDF format, stored on a network share or on internal storage.
- PPD Files: User-uploaded PPD driver files.
- All User Files: Shows PDF form files and any files generated with the print diagnostics feature (enabled in **Print Platform > Queue Config > Output**).

To upload a file to the PdfPlus, use the “Upload a new file” field.

Each page of the File Manager contains a search field which will display files whose name matches the entered search term.

Where relevant, listed files can be viewed, printed and deleted from the system using the icon buttons in the “Action” column. When an option to print is given, the output queue to which the files will be sent is chosen from the “Output queue” menu.

8.1. Form Files

The file browser shows a list of the files stored on the internal flash or hard disk of the PdfPlus. Note that not all files stored on the PdfPlus are displayed. Archived files and log files are not displayed here; the list is intended for user-supplied files such as PDF overlays and printer macros, and diagnostic files which may be generated by the PdfPlus job-processing software.

8.2. Log Files

Shows daily system logs stored either on a network share, or on internal storage. System log settings are configured in **System Settings > System Log**.

8.3. Archive Files

Shows files archived to a network share or internal storage. Choose the queue to list archived files for in the menu “List files for queue”. The number of files shown can be limited with the menu “Max files to display”.

8.4. PPD Files

Shows PPD files and gzipped PPD files which have been uploaded to the PdfPlus by the user.

8.5. All User Files

This page displays PDF form files and other user-uploaded files. Files generated with the print diagnostics feature (enabled in **Print Platform > Queue Config > Output**) appear here.

Print Diagnostics

For diagnostic purposes, you can have the print job saved before and after it is processed by the PdfPlus software. The input data will appear in the file list as prefilter_queue.bin and the output as postfilter_queue.bin.

This feature has the potential to fill the storage of the device and lead to failed print jobs, so care must be taken with its use. Because of this, the feature will always be switched off following a reboot of the PdfPlus.

Enable print queue diagnostics in the Output tab of **Print Platform > Queue Config**.

Chapter 9. Print Reports

The PdfPlus can provide an audit trail showing what time each job is received by the device, and what time it is sent on to the printer. Print reports are generated from analysis of system log data. The reports can be viewed via the web interface and emailed daily to a group of recipients.

Because the reports are generated from system log data, it is necessary to configure the storage of system logs either to a network share, or to local hard disk storage. Print reports cannot be generated from the volatile system log which is always available from RAM.

Refer to the section called "System logging" in Chapter 4 for instruction on configuring the system log settings. Ensure that **Store log messages to network share or local hard drive** is checked, as shown in Figure 9.1, so that the system log will be permanently stored. The system log will be rotated daily in this case, creating a job log file and a system log file each day, as described in the section called "System logging" in Chapter 4.

Figure 9.1. Syslog settings necessary for report generation

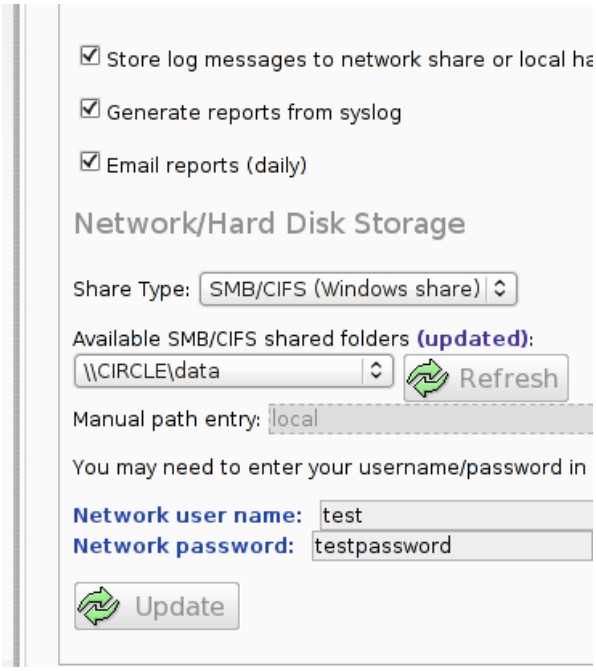

Checking **Generate reports from syslog** will ensure that reports are generated from the daily log files. Reports are generated each morning by processing the YYYY-MM-DD-wmlpp_job.log file. The report is output as YYYY-MM-DD-wmlpp_report.csv. Note that the report date will be one day earlier than the job log date. 2009-07-12-wmlpp_job.log contains the job logs for 2009-07-11 and is used to generate 2009-07-11-wmlpp_report.csv, which is the report for 2009-07-11.

If **Email reports (daily)** is checked, the YYYY-MM-DD-wmlpp_report.csv report will be emailed as an attachment each day, along with a summary of the day's printing. An example is shown in Example 9.1.

Example 9.1. Example email report

```
From:    wmlpp@wmltd.co.uk
To:      recipient1@example.com, recipient2@example.com
Subject: Report: 2009-07-02-wmlpp_report.csv (from wmlexample2)
Date:   Fri, 03 Jul 2009 06:04:33 +0100


WMLPP Printing Report: 2009-07-02-tempest_report.csv

Summary for wmlexample2:

Total jobs received:                      1043
Batch_A:                                  179
Batch_B:                                  303
raw_input:                                550
Unmatched jobs:                           11
Total jobs successfully transferred:      1029
Total jobs not yet transferred:           14
Total Megabytes transferred:              821.559
Total simplex sides received for printing: 29842
Total sheets if printed duplex:           14921

A detailed report is attached to this email.

Thank you for using the WML Print Platform.

Attachment: CSV document attachment (2009-07-02-tempest_report.csv)
```

The format of the report is comma-separated values, with each field enclosed in inverted commas and separated by a comma. The summary information is presented first, followed by a detailed listing of each job. An example is shown in Example 9.2.

Example 9.2. Example report attachment

```
"Summary of jobs sent by wmlpp1:"
"Total Jobs","Un-matched","Total MB","Total Sides Printed","Total Duplex\
 Pages","Total Completed","Total Not Yet Complete"
"13","13","0.173437","24","16","4","9"

"Detailed listing of jobs sent by wmlpp1:"
"Job Title","Bytes","ID","Simplex Sheets","Duplex Sheets","Received at \
wmlpp1","Complete Send to Print Device"
"qlpd_job","128","1","3","2","14/07/2009 10:21:26","14/07/2009 14:27:39"
```

9.1. Viewing reports in the web interface

To access reports in the web interface, log in and select **Print Reports**. Your browser will display a page similar to Figure 9.2

Figure 9.2. View of available reports

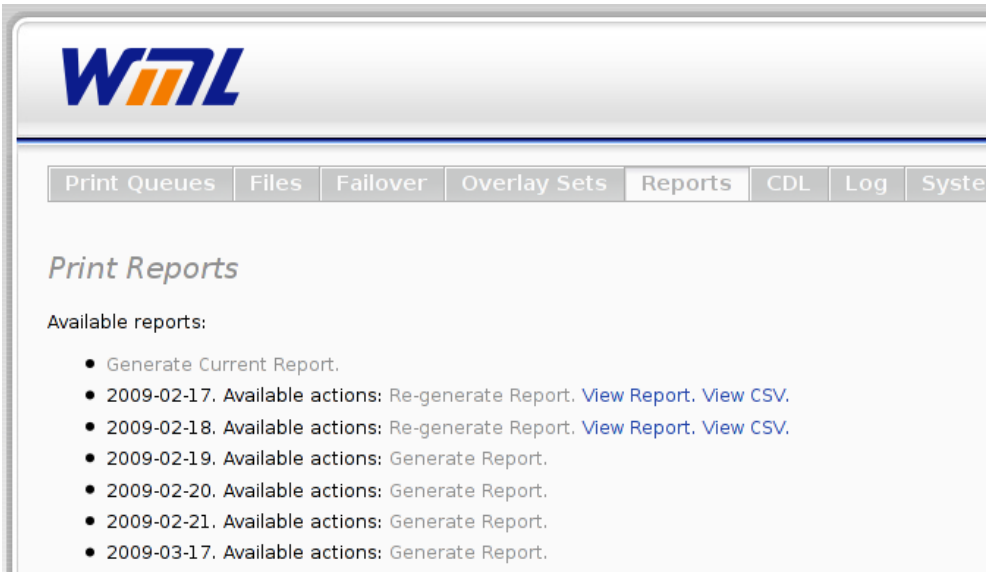

Available reports are shown in a list. The very first entry is **Generate Current Report**. This will generate an up-to-the-minute report based on the current print log. If you click the link, the print log will be processed and then two new links will appear. **View Report** and **View CSV**. Click on either to view the report.

For each dated job log file which has been created in the system log store, an entry is seen on the Print Reports page. Here, you have the option to generate, re-generate or view the report. If you generate a report, a second counter will count the time as the report is generated.

> **Note**
>
> For a production system printing hundreds or thousands of jobs per day, report generation may take as long as a few minutes to complete.

Figure 9.3. Viewing a report

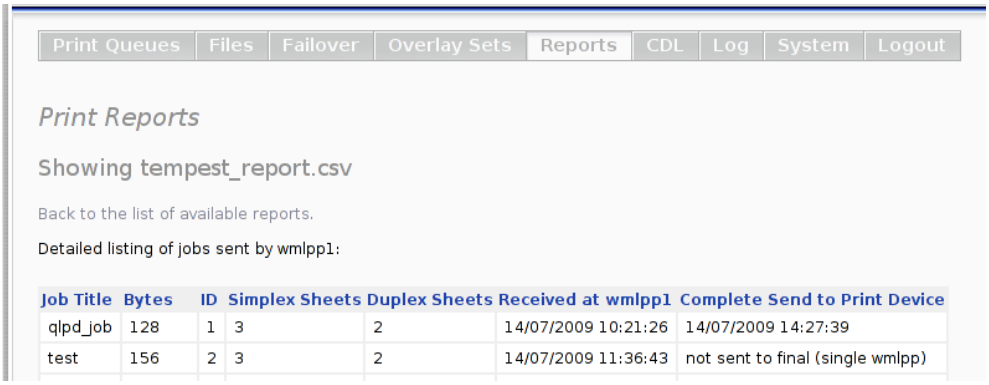

Figure 9.3 shows an example of what you will see if you click **View Report.** You can return to the list of reports by clicking the link **Back to the list of available reports.** The listing of jobs is presented as a table, containing the title, its size, the internal job ID and the number of pages, along with the time the job was received and when it was sent on to the printer.

Figure 9.4 shows the same example data, but on the CSV page, where the data is presented in comma-separated format.

Figure 9.4. Accessing report in CSV format in the web interface

9.2. Reports with a pair of DataGateway devices

The WML DataGateway is able to work in pairs, in which data is transferred from one DataGateway to the other, ensuring secure, audited, automated data transfer across a wide area link.

> **Note**
>
> This feature is not available on other WML products, such as PdfPlus and PrePrint.

The log messages from each DataGateway are stored on both systems. An SSL tunnel is used to transfer the messages in real-time. The ports used by the SSL tunnel are configured in **System Settings > System Logging**. **Send syslog messages to** should be set to protocol TCP/SSL and the suggested port to use is **5140**. **Receive network syslog messages** should have protocol set to TCP/SSL, and again it is suggested that port **5140** is used. This port will need to be opened on any firewalls which are configured between the DataGateways.

In addition to the real-time log messages, an SSL based secure copy process is used to transfer some files (especially logs recording barcode scans of documents). To enable this transfer, the sending DataGateway should be permitted to make a connection on TCP port **7522** to the receiving DataGateway.

Chapter 10. Kyocera Prescribe Emulation

The Prescribe emulation is offered as an optional module for PdfPlus. The emulation can be enabled by the use of a suitable key in **System > Licences**.

WML's Kyocera Prescribe Emulation converts Prescribe input data, including formatted text and simple PCL, directly into Portable Document Format (PDF).

The converted PDF files can be archived to a document store provided by a Windows or Unix network-share and finally processed for printing through the on-board printer drivers, which convert the document into PCL or PostScript.

The Kyocera Prescribe emulation is configured in the **Queue Config** section of the PdfPlus user interface in the tab **Kyo Settings**. Each print queue can be configured with unique Prescribe parameters.

The emulation parameters specified in **Kyo Settings** are effectively the “printer defaults”. The parameters should be configured to match the settings on the printer on which the data was previously being rendered.

The WML Kyocera Prescribe emulation follows the definition of the Prescribe page description language given in the Prescribe Technical Reference 4.51, available from Kyocera.

Figure 10.1. Prescribe emulation configuration interface

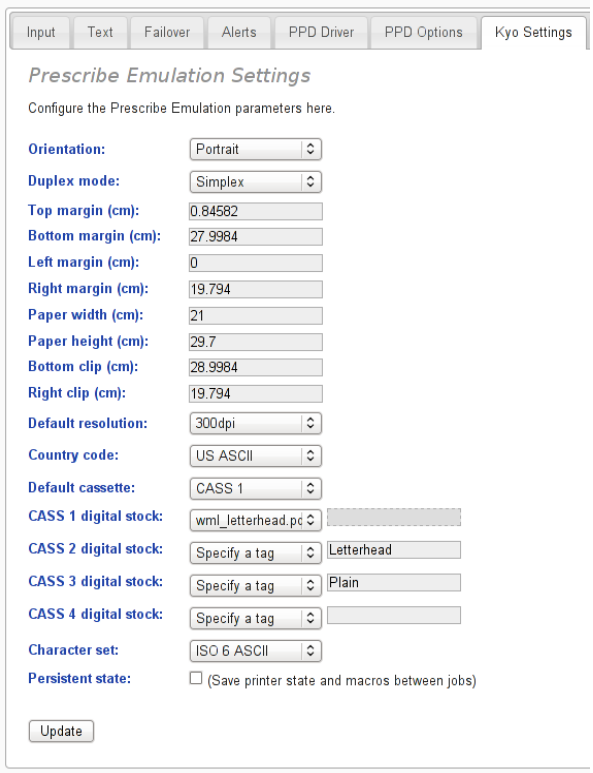

10.1. Basic Defaults

The basic parameters, *orientation*, *duplex mode*, *default margins*, *paper width* and *height* and *clipping* parameters, will be familiar to all users and function according to the Prescribe specification.

The *country code* and *character set* are used to select the appropriate internal character set and glyphs according to the Prescribe specification, where possible.

10.2. Paper Cassettes

The *default cassette* is the paper cassette which the emulation will use if the cassette is otherwise unspecified (with a CASS or PSRC command).

The PdfPlus has two ways to apply paper stock to your printing. It can either apply *digital paper stock* to the output, using a PDF file as the background to the print, or it can embed *PostScript paper selection tags* into the output data.

> **Note**
>
> When used with a PCL printer, only digital letter stock can be used to emulate the CASS/PSRC Prescribe commands.

Using Digital Paper Stock

Digital stock is especially useful when building electronically distributable documents. Digital stock can be used to add a letterhead or the company logo to documents which were previously printed on pre-printed stock.

To set digital stock for a particular cassette, you must first upload the PDF of the stock onto the PdfPlus appliance. Got to **File Manager > Form Files**. Use the "Upload a new file" field to upload the PDF. Newly uploaded PDF files will now appear in the menus in **Kyo Settings**. To set your newly uploaded PDF file as the stock for Cassette 1, select it from the menu for "CASS 1 digital stock".

Using PostScript Paper Selection Tags

To have a paper selection tag embedded in PostScript output for Cassette 1, select "Specify a tag" from the menu for "CASS 1 digital stock". Type the exact text of the tag into the adjacent text field.

> **Note**
>
> Common text tags are "Plain" and "Letterhead" but be aware that these are often case sensitive and each printer will have its own tags. In some cases, the tags can be user-specified in the printer or RIP environment.

10.3. Persistent State

If "Persistent state" is checked, then the WML emulation will save the effective printer state between print jobs. This ensures that Prescribe macros are stored between jobs. It will also store the current paper cassette value, page margins and so on.

Storing and re-reading the printer state between print jobs has a small performance cost.

If you need your Prescribe macros to be stored between individually sent print batches, then check "Persistent state", otherwise leave it un-checked.

Chapter 11. Appliance Manager

Appliance Manager is an optional module for PdfPlus which allows you to manage a group of WML appliances. The module is enabled (with a licence key) on a single PdfPlus which becomes the master appliance manager on the network.

Appliance Manager has the following features:

- Catalog individual appliances and create appliance groups in a user-friendly web interface.
- View realtime status information for all appliances.
- Push appliance settings out to an appliance or a group of appliances.
- Apply firmware upgrades to an individual appliance or to a group.
- Push out a reboot or shutdown command to an appliance or group.
- Review managed appliances in a CSV-formatted report.

A time window can be set for all actions to ensure that they are carried out during off-peak hours.

All communication channels between appliance manager and client devices are 256-bit SSL encrypted ensuring that Appliance Manager can be used safely across a public wide-area network.

11.1. The Appliance Manager "push"

Appliance Manager provides the concept of a **push**. A push is the specification of a **configuration**, along with an appliance or group of appliances and a time window during which the push should be applied.

A configuration may be a **settings configuration**, a **firmware upgrade** or a **command**.

A settings configuration is a **backup.tar** file, as generated within **System Settings -> Backups**. It may contain a set of templates or print queues. The settings which can be incorporated into the configuration are listed in the section called "System settings".

A firmware upgrade file is a standard WML upgrade file, as provided for an individual appliance. A typical firmware upgrade file might have the filename **wmlpp-dg-kvm-sata-tree5-20110125_1142.img**.

Commands which can be distributed with a push currently include only the **reboot** and **shutdown** commands.

A managed appliance will periodically query the appliance manager to see if there are any pending pushes which it should apply. The manager then transfers the data (upgrade image or backup file) and waits for the appliance to apply the update and confirm success.

The main activity for the administrator using Appliance Manager is to create and manage pushes.

11.2. The user interface

Appliance Manager has three section to its interface:

- **Appliances and Groups**: The main list of appliances and appliance groups;
- **Stored Configurations**: A page to manage uploaded settings configurations (backup.tar files from individual appliances);
- **Stored Firmware**: A page to manage uploaded firmware upgrade files.

11.3. Appliances and groups

This section covers the most important part of the Appliance Manager user interface; **Appliances and Groups**.

Figure 11.1. The Appliances and Groups section of the Appliance Manager.

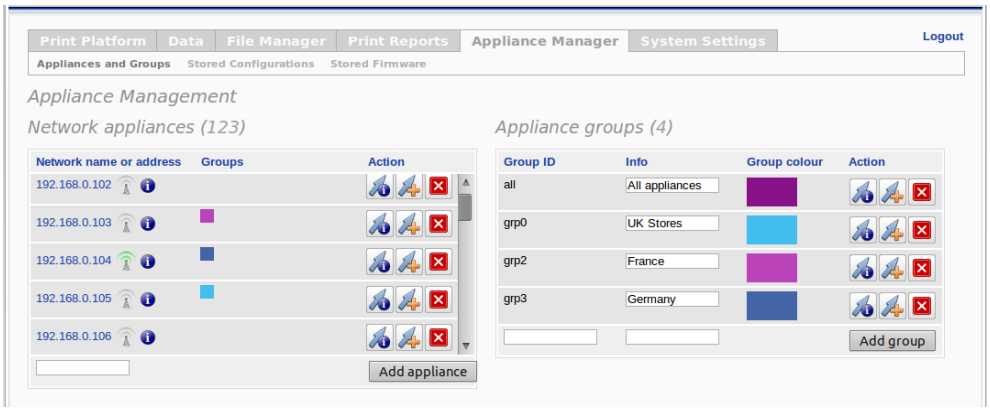

Figure 11.1 shows the **Appliances and Groups** page. To the left is a table headed **Network appliances** listing each appliance being managed. This list is created by the administrator, either by adding entries with the **Add Appliance** row of the table, or by uploading a CSV file containing a list of appliance addresses.

> **Note**
>
> The number of appliances and groups under management is shown in brackets after the table heading.

The **Network appliances** table displays the appliance address (IP address or network name), its online/offline status ( ), an appliance status button ( ), group memberships (denoted by coloured blocks) and a set of action buttons (, and).

> **Note**
>
> The "online status" icon will be green ( ) if the managed appliance is online and successfully sending status messages to the appliance manager.

> If the managed appliance is not sending status messages (even if it is otherwise accessible, for example by ping or the web-interface) then the icon will appear grey ( ).

The **Appliance groups** table to the right allows the administrator to create groups. A group is given a name (**Group ID**) and a short description (**Info**). It receives an auto-generated colour which is diplayed as a large block (**Group colour**). The action buttons in the **Appliance groups** table have the same meaning as those in the **Network appliances** list.

To make an appliance a member of a group, drag the appliance row over and drop it on the group row. The group colour will appear as a small block in the **Groups** column on the **Network appliances** list. To remove an appliance from the group, click on the group block (when you hover over the group block, a white cross will appear).

Populating the appliance list

Figure 11.2. Upload an externally generated list of appliances, or access an appliance report.

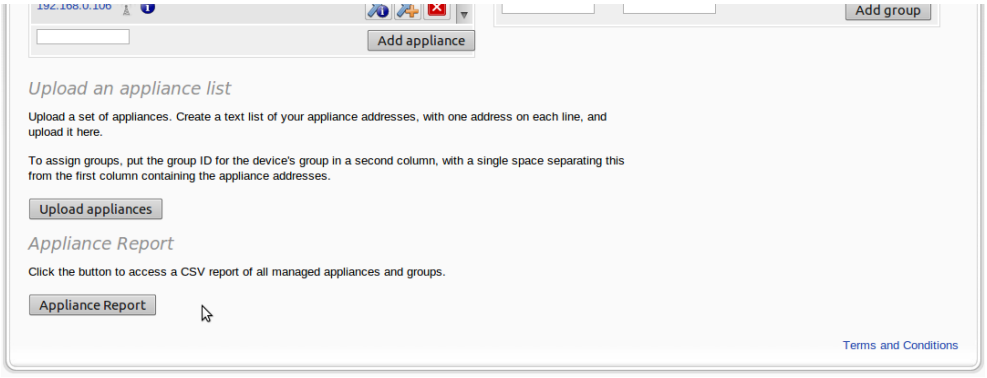

Towards the bottom of the Appliances and Groups page, there are two additional sections, shown in Figure 11.2. **Upload an appliance list** provides an interface to use a text file to populate the **Network appliances** list. **Appliance Report** generates a full report of the system and is covered later in this section.

This list of addresses can be a single column of IP addresses or network names as follows in Example 11.1:

Example 11.1. Appliance list (with no groups specified)

```
192.168.0.1
192.168.0.2
192.168.0.3
appliance1.local.lan
appliance2.local.lan
appliance3.local.lan
```

Addresses on this list are added to the existing list of network appliances without first deleting any pre-existing entries. If any of the appliances are already listed then duplicates are *not* added.

Example 11.2. Appliance list with group specified in column two

```
192.168.0.1 London
192.168.0.2 London
192.168.0.3 London
appliance1.local.lan London
appliance2.local.lan Reading
appliance3.local.lan Reading
```

To set group membership when uploading the list of appliances, include a second column of the appliance group, as shown in Example 11.2. The columns should be separated with a space. Column two is the group ID and should contain no spaces.

> **Note**
>
> If the group specified in the uploaded file does not exist, it will be created.
>
> If any groups *are* created, it will be necessary to re-load the **Appliances and Groups** page to display them, as they are not automatically added to the **Appliance groups** table.

Appliance status

To view the status of a managed appliance, click on the status icon ( ). The status dialog will be displayed, as seen in Figure 11.3.

> **Note**
>
> It will take up to five minutes for the system status to appear for a newly managed appliance.

Figure 11.3. Viewing the status of a managed appliance.

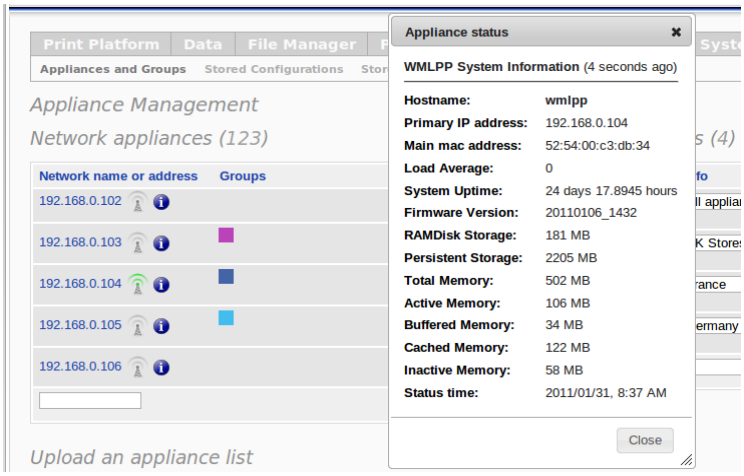

Most of the status fields displayed in **Appliance status** are self-explanatory, but note:

- The **Status time** displayed is the time as reported by the clock on the managed appliance.
- Of the memory fields, the most important is **Active memory**, which should be less than **Total memory** by a sensible margin. (The memory data is extracted from the Linux kernel's /proc/meminfo.)

When using Appliance Manager, it is important to ensure that all appliances have the same system time. The best way to achieve this is to make use of the NTP feature (See the section called “System date and time”).

> **Note**
>
> If the system time on the managed appliances differs by a significant amount from the appliance manager, then the SSL-encrypted communication between appliance and manager will fail. (SSL certificates do not validate if they appear to be too old, or exist in the future).

Appliance report

To view a full report, use the **Appliance Report** button. This will generate a comma separated value-formatted report containing:

- All managed appliances, with status
- Groups with lists of appliance members
- Full listing of Pushes, both complete and incomplete, including appliances to which the pushes apply.

Figure 11.4. Viewing the full report of managed appliances and groups.

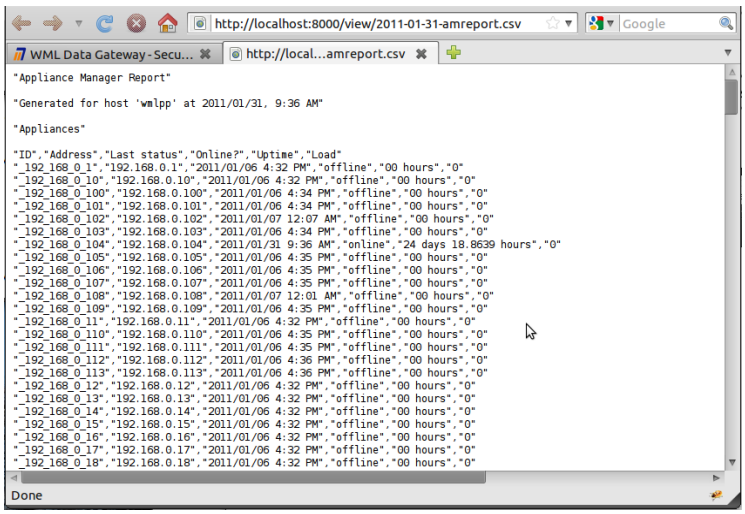

The report will pop-up in a new tab or window in your browser (as in Figure 11.4). Save this file and then import it into a spreadsheet to access and manipulate the information.

> **Note**
>
> You will need to set your web browser to allow pop-up windows for the appliance manager. If you do not do this, you will probably see a message stating that the browser blocked a pop-up when you try to create the appliance report.

Managing pushes

Figure 11.5. Creating a new "push", and choosing a configuration to push out to the appliance.

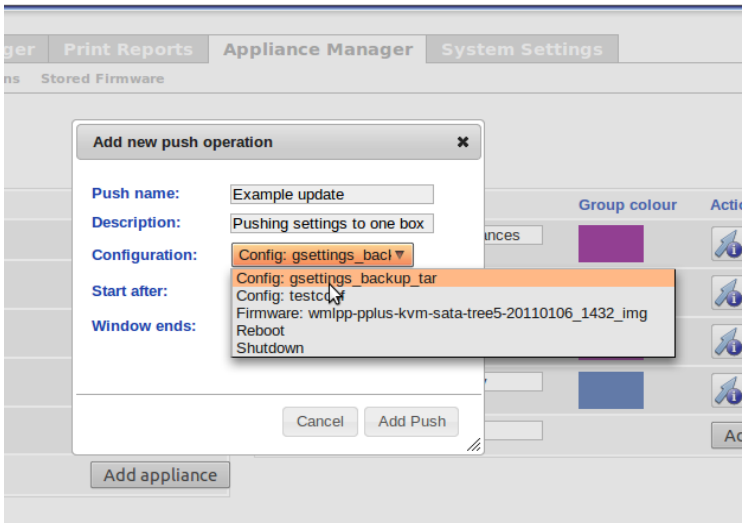

Use the add push icon ( ) to create a new push for an appliance or group. A dialog will pop up for the push details, as shown in Figure 11.5. Enter a name for the push, a short description and choose the configuration.

Figure 11.5 shows the configuration menu displaying five options; two example settings configurations, a single firmware upgrade file and the standard commands **Reboot** and **Shutdown**. The number of options in this menu depends on the number of configurations and firmware images which have been uploaded on the **Stored Configurations** and **Stored Firmware** pages.

If the push should be applied as soon as possible, choose **ASAP** in **Start after** (Figure 11.6). Otherwise, choose a start time and a time when the push window ends. The system will attempt to apply the push during that window. Unfortunately, it's not possible to specify a particular date for a push.

Figure 11.6. Choosing a time window for the new push.

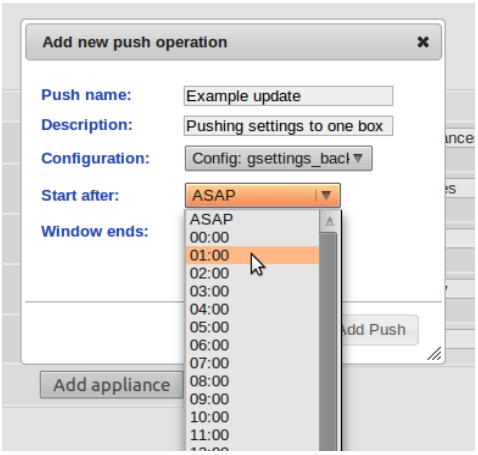

If this is a push for a group, then it will include all of the appliances in the group *at the time the push is created*. If an appliance is added to the group *after* the creation of the push, it will *not participate* in the push. Likewise, if an appliance is removed from the group after the push was created, it *will* still have the push applied (unless the appliance itself has been deleted).

Figure 11.7. The list of appliance "pushes" for the appliance "_192_168_0_104".

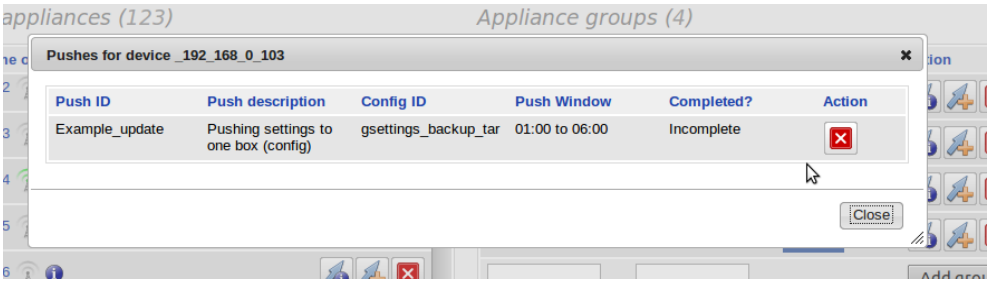

Once the push has been created, it can be viewed with the push query icon ( ). This will bring up a dialog box as shown in Figure 11.7. This is a simple table showing the pushes for the appliance or group. The ID of the push is shown, along with its description, the config which is to be pushed out, the time window and whether or not it is complete. If the push *is* complete, the time of completion is shown.

> **Note**
>
> The appliance ID in Figure 11.7 is _192_168_0_104. This is generated from the appliance address, by conditioning the address into a tag which can be used within Appliance Manager's internal XML configuration files, hence the '_' characters replacing '.' characters.

If you hover your mouse over the **Push ID** field, the appliance members of the push are shown. This is useful when reviewing group pushes. Likewise, if you hover over **Completed?** a list of the appliances which have applied the push is shown as well as those which have yet to complete.

Completed pushes are displayed until the administrator decides to delete them. A push can be deleted ( ) whether it is complete or not. If an incomplete push is deleted, those appliances which have not applied the push will stop attempting to do so.

11.4. Stored configurations

This page (Figure 11.8) provides an interface to upload **backup.tar** files for distribution to managed appliances.

Figure 11.8. An interface to upload settings configurations (backup.tar files).

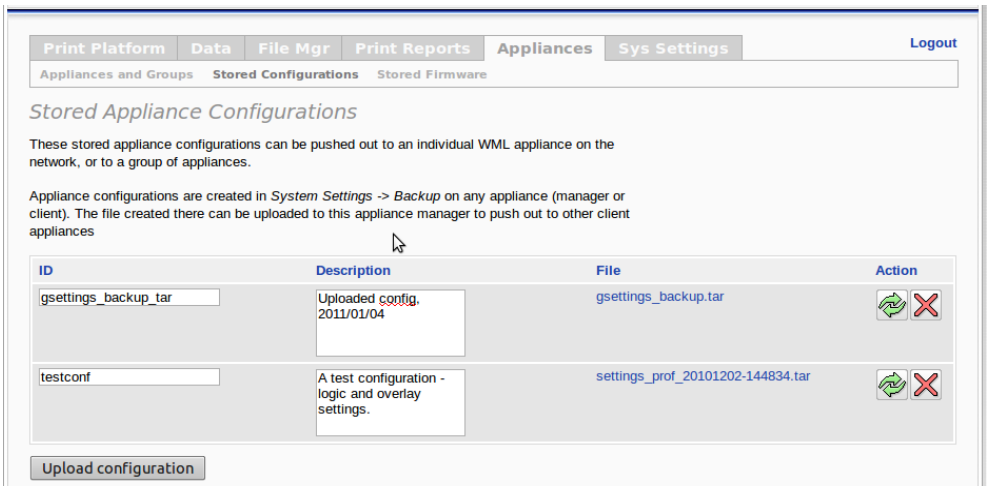

Typically, an administrator would configure a separate appliance with suitable settings. This could be one of the managed appliances, but it does not have to be. He would then save the configuration for the appliance in **System Settings -> Backup**, choosing only the "section" of the settings required. For example, he might save off the templates and logic settings, so that these could be pushed out to the managed appliances without affecting their system settings and print queues.

He uploads the backup.tar file to this interface where it is given an ID based on the file name and a description based on the date. The configuration will now appear in the drop-down menus of the **Add new push operation** dialog in the **Appliances and groups** page.

> **Caution**
>
> The interface *will* allow you to delete a configuration which is a member of a pending push. If you do this, then the push will not complete.

The **Action** column provides update and delete buttons to update changes to the **ID** or **Description**, or to delete the entry.

11.5. Stored firmware

Figure 11.9. Upload firmware upgrade files here, ready to be pushed out to managed appliances.

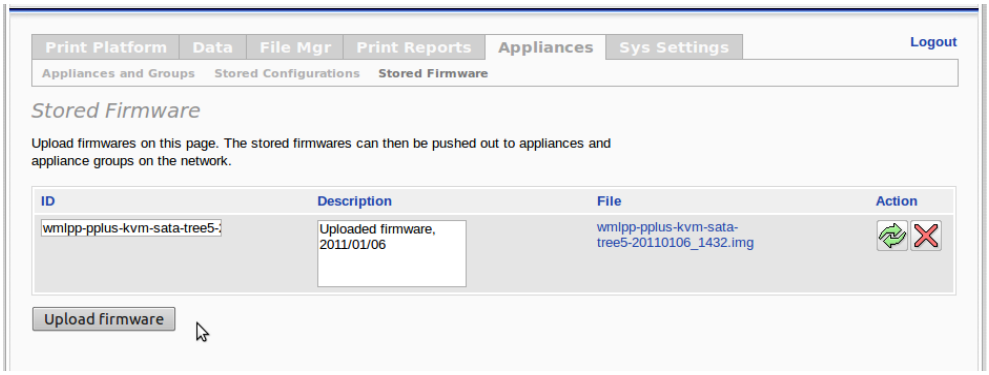

This page (Figure 11.8) provides an interface similar to that in **Stored Configurations**. This is for uploading and managing firmware upgrade files.

The administrator uploads the same firmware upgrade image which he would use to upgrade a single box. Again, it is given an ID based on the file name and a description based on the date. Once uploaded, it is available to be pushed out to upgrade managed appliances.

> **Caution**
>
> The interface *will* allow you to delete a firmware which is a member of a pending push. If you do this, then the push will never complete.

As in **Stored Configurations**, the **Action** column provides update and delete buttons to update changes to the **ID** or **Description**, or to delete the firmware file.

11.6. Appliance Manager for Disaster Recovery

The Appliance Manager system is used to provide a Disaster Recovery (DR) backup system. This allows the firmware, config and data on one PdfPlus to be mirrored to a second PdfPlus.

To set up a DR pair, designate the master system. Appliance Manager should be licenced on the master system. On the master system, add a single network appliance in **Appliances > Appliances and Groups**. Enter the correct IP address for the DR backup system and press **Add appliance**. Check the **DR** checkbox to designate this as a DR backup system.

All configuration settings and firmware upgrades will now be transferred to the DR backup system. Settings and data (including the contents of the Document Content Control

database) are transferred twice daily. Firmware upgrades are pushed from the main system to the backup system shortly after the main system is upgraded.

> **Warning**
>
> The DR backup system should be running on the same infrastructure as the main system. Either both should be KVM based (such as those which are supplied by WML as hardware appliances) OR both should be VMware based. If one system is KVM based, and the second is VMware based, firmware upgrades applied to the main system will not propagate to the DR backup system.

The DR backup will hold its existing IP address settings until it is "rebooted as a DR main system". This is done by accessing **System Settings > Status/Control** on the DR backup system. On this system there should be a reboot icon labelled **Disaster Recovery Reboot**. If this is pressed, the DR backup system will reboot and take on the personality of the DR main system, including its IP address settings.

> **Warning**
>
> It is very important to ensure that the main system is offline when the DR back-up is brought up with **Disaster Recovery Reboot**, otherwise IP address conflicts will ensue.

11.7. Appliance Manager Networking

This section describes the ports used by Appliance Manager to communicate with managed devices.

The Appliance Manager **listens** on **UDP port 1554** for incoming connections from managed devices so that these devices can announce their presence. This allows newly added managed devices to be automatically listed in **Appliances > Appliances and Groups**.

The Appliance Manager has a server **listening** on **TCP port 1555** for status messages sent from managed devices. This channel is SSL encrypted.

The Appliance Manager **makes connections** to **TCP port 1556** of managed devices to send firmware, config and command pushes. This channel is SSL encrypted.

Finally, Appliance Managers which have managed DR partners **make connections** to **TCP port 22** on the DR partner(s) to transfer configuration settings. All communications on this channel are SSL encrypted.

Chapter 12. WML Secure Forms

WML Secure Forms is a module for PdfPlus which extends the overlay template editor (see "Overlay template manager"). It presents a view of an overlay template canvas which can be filled in by a customer or by a sales assistant working with the customer.

Secure Forms is enabled by entering a suitable licence key in **System > Licences**.

Typically, an overlay template is set up with a PDF form as the background, providing graphical content, and a set of **variable data** elements. In the Secure Forms interface, each variable element is presented as a text-entry field.

When the form has been completed, the sales assistant selects the **Print** option and the document is printed and archived.

> **Note**
>
> Why *Secure* Forms? Because the graphical content of the form is *not held on the client device*.
>
> For example, Imagine a retail store is using Secure Forms with tablet devices. The sales assistant has customers fill their details into the form using the tablet.
>
> If the tablet is stolen, the ability to reproduce the form (which is likely to be a financial document) is not taken with the tablet; it's still locked in the store's server room.

Secure Forms is implemented using web technology. It's simply a web page. Like the rest of the PdfPlus it makes extensive use of Javascript, so client devices must conform to the minimum browser requirements.

Figure 12.1. A template ready to be presented as a Secure Form

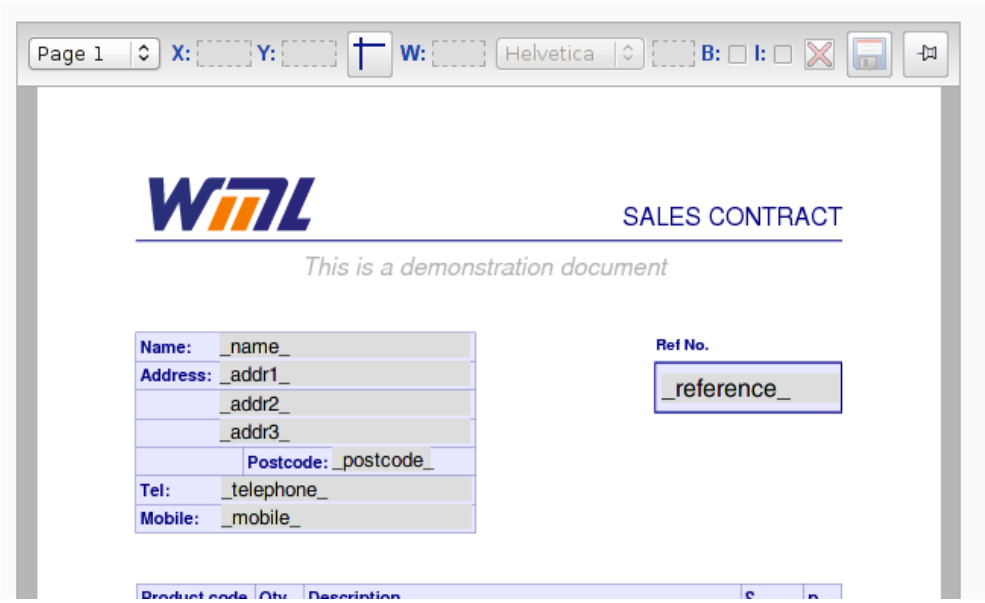

Figure 12.1 shows a sample template in the overlay template editor. It consists of a PDF form background on which numerous variable data elements are placed, with key names such as name, addr1, addr2 and so on.

Figure 12.2. The Secure Form login

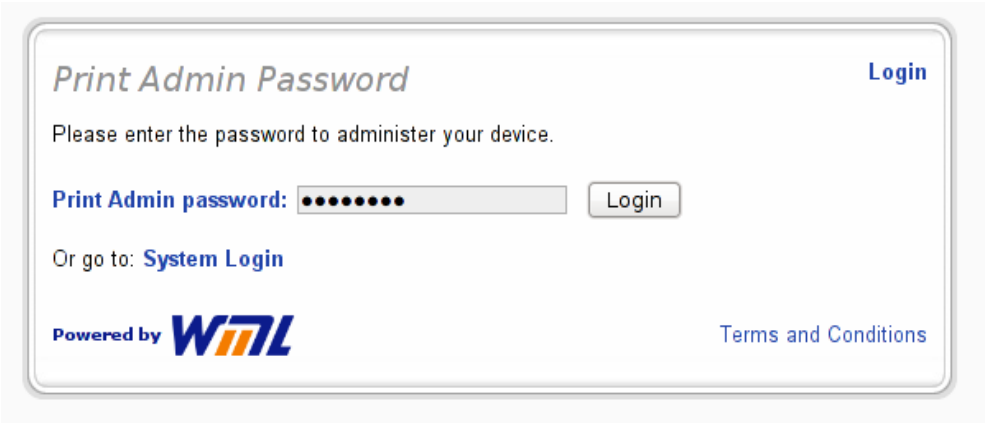

To access the Secure Forms interface, log out of any PdfPlus system admin session you have. Redirect your browser to http://PdfPlus.address/formportal.cgi. You will see a login page as shown in Figure 12.2. The password is, by default, the same as for the system admin interface (i.e. “blankdoc”).

Figure 12.3. The Secure Form interface

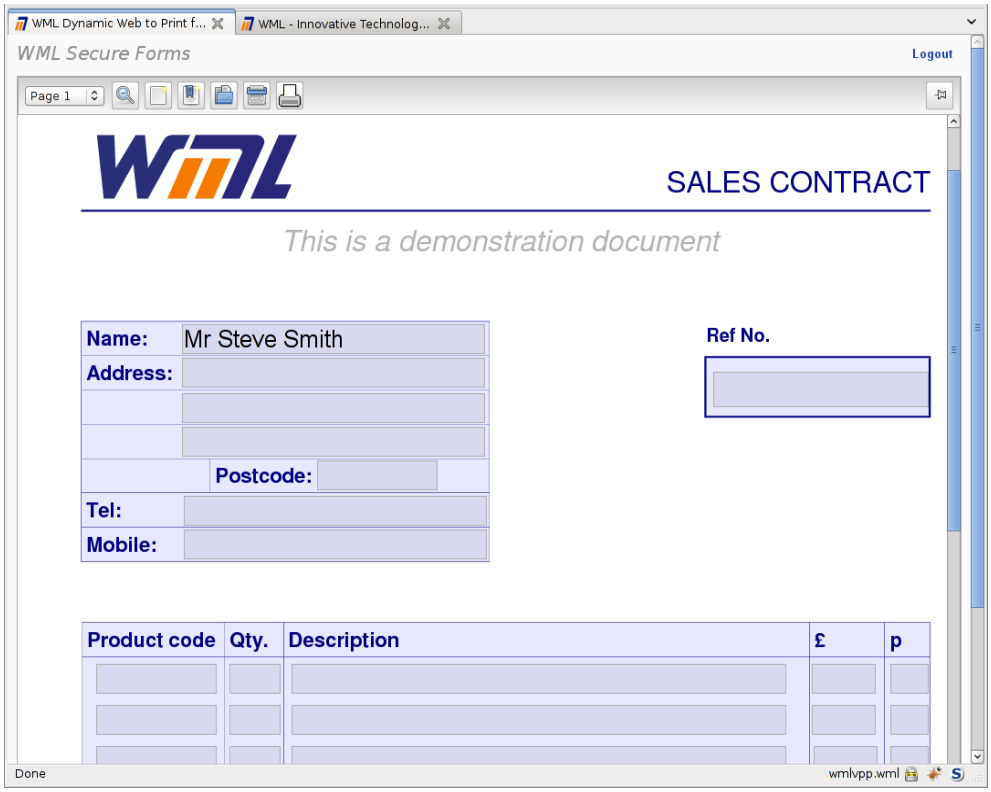

Figure 12.3 shows the Secure Forms interface which would be presented to the sales assistance and the customer. It's a very simple interface with the form laid out as it would print and a top toolbar.

The form is initially presented to fit the page-width of the browser window. As soon as one of the text fields is selected, the form “zooms in”. Click outside any text field and the view zooms out again. There is a separate button to action the zoom on the toolbar ( ).

Every time an element is filled in, the semi-complete form is saved on the PdfPlus. The saved form is automatically named “Untitled data”, “Untitled data(1)” etc. As soon as the form is printed, the saved form data is erased.

Figure 12.4. Bookmarking an incomplete form

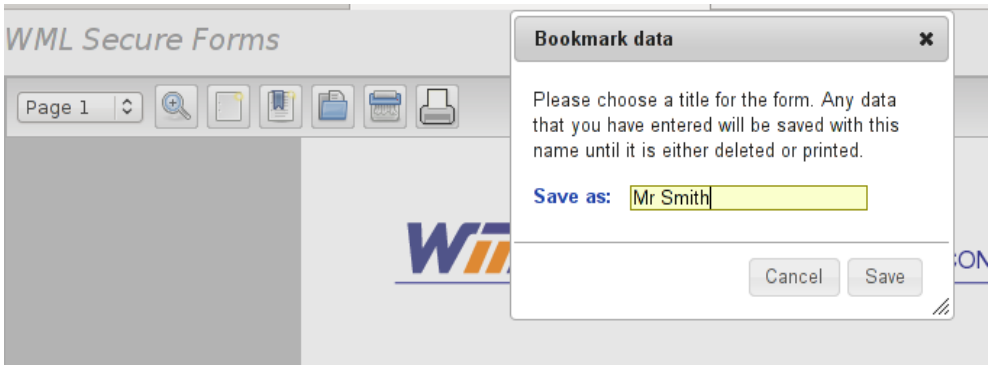

It's possible to name the current form data using the bookmark button ( ). To load a previously bookmarked form, use the open button ( ).

If a form is incomplete, but is no longer required (i.e. it will not be printed), delete the data using the shred button ( ). To start a new named form, press the new button ( ) and provide a name.

Figure 12.5. Printing the form

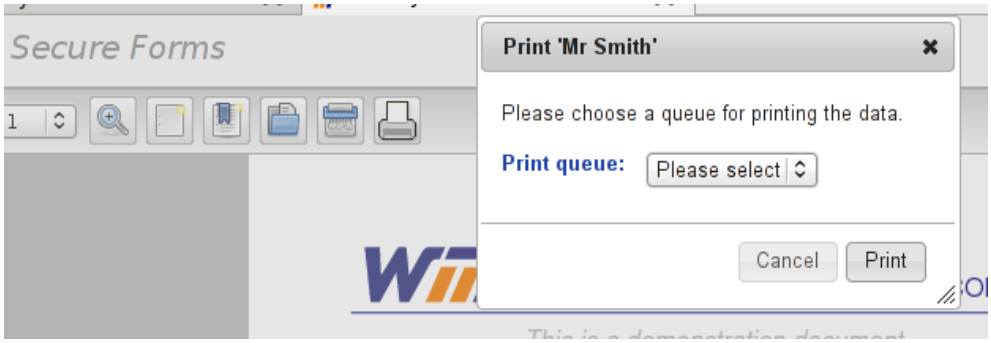

When the form is complete, press the print button ( ). A dialog will open allowing you to select the destination printer (this is a list of all the queues which are configured on the PdfPlus). Click **Print** and the form will print to the selected printer.

> **Note**
>
> Most installations of Secure Forms will require customisation to add database interactivity, either to store customer details or to automatically retrieve information for existing customers.
>
> Please contact WML to discuss your individual requirements for a Secure Forms installation.