Pacific Knowledge Systems

RippleDown Deployment Guide
Version 8.1

This document outlines the deployment of the RippleDown application

Copyright Notice

The information provided in this User's Guide is subject to change without notice and is not a commitment by Pacific Knowledge Systems Pty Ltd. The software described in this User's Guide is provided under a license or non-disclosure agreement. It is unlawful to copy this software except as allowed in the agreement.

No part of this User's Guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information retrieval systems, for any purpose other than for the purchaser's personal use, without the written consent of Pacific Knowledge Systems Pty Ltd.

Reproduction or disassembly of embodied programs or databases that make up the software is prohibited.

© Copyright Pacific Knowledge Systems Pty Ltd, 2019

All Rights Reserved.

Intended Use

The PKS Software, when used as a medical device, is intended to be used purely as a decision support system that provides complementary reports for patient data to qualified individuals based upon established rules set up by relevant trained customer domain experts (hereafter referred to as "domain experts").

The PKS Software generates and collates comments into a patient-centric report or workflow action based on rules created and maintained by domain experts. The PKS software presents all first-of-a-kind reports for review, modification (if applicable) and approval by the domain expert prior to release to clinicians or other individuals. Beyond this, the domain expert may automate the release of none, some or all identical reports. However, since a report may be generated and automatically approved for a case for which the domain expert has not previously considered, it is intended that the domain expert regularly review a representative sample of all output types.

It is intended that the clinicians receiving reports against patient results will consider the report in conjunction with all current and previous patient history and apply their own judgement when determining patient management. It is intended that the clinician does not rely on the existence of a report for the management of a patient.

The PKS software in itself does not provide automated diagnosis or treatment-making functions or have the capability to control, in any way, the performance of a device or to treat or diagnose any disease.

It is the responsibility of the licensee to use the product in accordance with its intended use. In support of this, it is recommended domain experts include a statement consistent with the following on all reports.

“This report has been generated using clinical decision support software. This report is intended to provide adjunctive information only and should not replace clinical judgement.

Disclaimer

Pacific Knowledge Systems Pty Ltd makes no warranties, either express or implied, regarding these computer software packages, or their fitness for any particular purpose other than warranty provisions embodied in any agreement or purchase contract.

Acknowledgments

Microsoft and Windows are registered trademarks of Microsoft Corporation.

Manufactured By

Pacific Knowledge Systems: http://support.pks.com.au/product-register/

RippleDown is distributed by Abbott Laboratories as AlinIQ CDS

RippleDown is distributed by Philips Healthcare as LABSOSYS CDS

Table of Contents

On-site Installation Components 5

Supported Interfaces 6

Integration with the Clinical Data Source 6

Platform Requirements 8

Security Recommendations 10

Remote Application Management 11

Log File Transmission 12

Decimal separator 12

A typical deployment schedule 13

On-site Installation Components

For on-site installations (as distinct from a cloud implementation), the following two installers are provided, and these can be run in unattended mode if required. Both Windows and Linux installers are available. Detailed installation and upgrade guides are provided.

Installer	Application	Summary Description
RippleDownServer	RippleDown Server	Communication interfaces to Clinical Data Source or instrument Interpretation engine Database management
RippleDownClient	Knowledge Builder	Facilities for senior staff to create and maintain their clinical or audit Knowledge Bases.
	Validator	Facilities for scientific or clinical staff to review and approve interpretations and correct them if necessary prior to release.
	Auditor	Facilities for data entry staff to review and correct orders which have been flagged by a Knowledge Base as possibly containing errors.
	Administrator	Facilities for IT staff to manage: Communications interface between the Clinical Data Source and RippleDown Project-based validation options User accounts Language translations Event log Database backups
	Translator	Facilities to allow reports to be generated in several languages from the one Knowledge Base

All 3rd party software used in RippleDown is loaded automatically by the installers and is licenced for your unlimited use.

The RippleDown server uses a high-performance, no-SQL, embedded database management system. No database management are required for its installation or maintenance. RippleDown databases are automatically compacted and copied daily during housekeeping to a directory where they can be backed up via your existing backup procedures.

RippleDown does not add or change any system environment settings or registry settings, apart from adding the installation entry.

Supported Interfaces

Protocol	REST
Type	Omnilab AMS

Integration with the Clinical Data Source

The following shows a typical integration of RippleDown Expert with the clinical data source (e.g. Hospital Information System (HIS), Laboratory Information System (LIS)) deployed at a client site (The integration of RippleDown Auditor is similar).

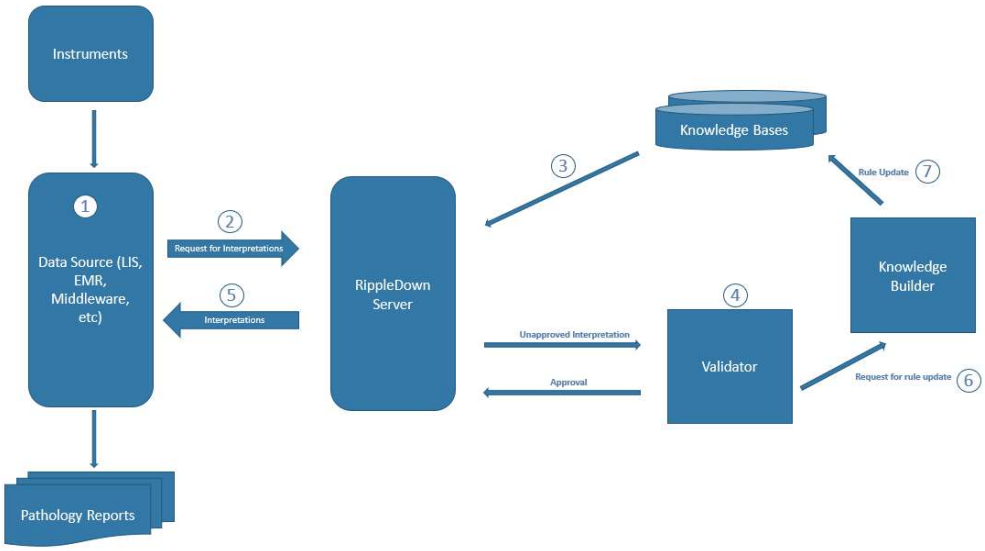

Step	Description
1	RippleDown is "spliced" into the Clinical Data Source at the point in time when test results from the analysers have been filed and verified.
2	The Clinical Data Source generates a “Request For Interpretation” message consisting of all the information required for the interpretation. This will typically include current and previous test results, patient demographic and visit data, referring and copy doctors, and clinical notes. RippleDown does not maintain its own database of patient data to retrieve past results, but rather relies on the Clinical Data Source to send the most up to date information for each transaction. If the Clinical Data Source is unable or unwilling to create all or part of this message, RippleDown may be able to retrieve this information from the Clinical Data Source or a Patient Data Repository with each transaction.
3	The RippleDown Server selects the appropriate Knowledge Base for this message and generates the corresponding interpretation. As well as the patient-specific report and workflow actions where relevant, the interpretation will include a flag indicating whether this interpretation is to be autovalidated or manually reviewed. Like all other Knowledge Base outputs, the autovalidation flags are generated by rules defined by domain experts/clinicians.
4	Interpretations that are not auto-validated are queued to the Validator where Validator users review them and make any required corrections prior to release. Particular types of interpretations can be queued to the most appropriate Validator queue. Several Validator users can be reviewing in parallel. Other validation workflows are supported, including where validation is performed within the Clinical Data Source rather than within RippleDown.
5	The validated “Interpretation message” is sent back to the Clinical Data Source. If the message contains workflow actions, such as reflexing a test, generating an alert, changing a billing status or other database item, these are done by the Clinical Data Source at this time. The interpretive report contained in the message is added to the patient results, incorporated into the report and delivered to the referrer..
6	If a Validator user made a change to the interpretation before approval, the case can also be queued to the Knowledge Builder for a rule update. This process allows your clinical expert maintaining that Knowledge Base to update it according to your Laboratory’s commenting preferences or other guidelines.
7	When a clinical expert updates the Knowledge Base with the new rule, subsequent cases that are similar to the original will have the amended interpretation applied. Corrections made by Validator users are therefore an important part of the knowledge building

Step	Description
	process, resulting in the continuous refinement and greater sophistication of the Knowledge Bases. RippleDown has a built-in conflict checking mechanism that activates during new rule implementation or rule update. That means that updates can be applied immediately within the production RippleDown system, without requiring a separate off-line testing phase.

Platform Requirements

An implementation consists of the RippleDown server and RippleDown thick client applications installed on user workstations.

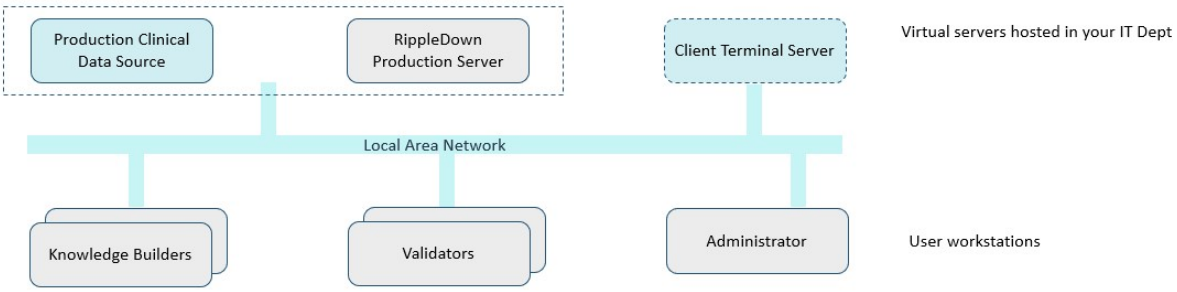

RippleDown has been baselined against the following dedicated infrastructure:

- Intel Xeon E5-2676 v3 @ 2.40GHz 2x Core
- 8GB RAM

Testing concluded that RippleDown will run smoothly with the following metrics while also leaving headroom for additional transactions.

- Support 10 concurrent UI users
- Process 10,000 automatic approvals per hour
- Process 500 manual approvals per hour
- Process 20 case searches per hour
- Sustained activity for an 8-hour period without any performance degradation

Case Search
Conducting searches on very large archive case lists can have a significant overhead on CPU utilization. It can affect the response time of case processing and user interaction with the thick client applications. It is recommended that the archive case list be configured to hold a maximum of 30 days of cases for all projects.

If an archive case list needs to be increased, it is recommended searching of larger archive case lists is used appropriately and / or outside of peak processing hours.

Searching a large archive case list may also cause the corresponding project's database to grow significantly and quickly which in turn may impact the Housekeeping batch job duration.

If a large archive case list is required for a project, it is recommended you seek guidance from support@pks.com.au prior to implementing the configuration change.

It is recommended for security and performance reasons that the RippleDown server application be installed on a dedicated server platform.

The RippleDown server application may be installed on the same platform as the Clinical Data Source, but the following additional requirements apply:

The security of the RippleDown Server or Clinical Data Source application could be compromised (e.g. the platform would need to be accessed by both the Clinical Data Source and RippleDown administrator users) unless additional security measures are implemented to mitigate this risk.

The resource availability of the co-located platform (e.g. CPU, RAM, disk) would need to be continually monitored and assessed to ensure that both applications have the required resources,

The co-located configuration would need to be validated against the testing volume profile in a test environment.

The minimum server specifications for a RippleDown server as outlined below.

Platform	Site Profile	Requirements
Production RippleDown Server	• Number of clinical databases <= 15 • Number of interpretations < 10,000 / day • "Days to keep" archive cases <=30 days	The performance of RippleDown has been baselined against the following dedicated infrastructure configuration. When co-locating on the same server as the Clinical Data Source, access to the equivalent resources will increase the likelihood of optimal performance of RippleDown. • Two 64-bit CPUs • 8 GB RAM

Platform	Site Profile	Requirements
		• 4 GB VRAM assigned to the RippleDown JVM (installation default setting) • 100 GB free disk, backed up daily • Windows Server 2008+ or Linux that can support Java JRE 8u202 • Access to SMTP server (for sending email alerts) • VPN access (for remote access by PKS Support)
User Workstations (for Thick Client applications)		• Single 32-bit or 64-bit CPU • Windows that can support Java JRE 8u202 • 10 GB free disk for installation • TCP/IP connection to the RippleDown Server • Internet access and browser to access product online help

Terminal Services
Thick client applications can also be installed on a Terminal Server or Citrix desktop solution, which may simplify management of the thick clients.

Security Recommendations

All RippleDown data (including backups, user credentials, report comments and rules) reside in unencrypted form on the server. PKS recommends you implement security measures appropriate to the needs of your site. Whilst highly unlikely, without adequate protection, reports and other outputs generated from RippleDown may be inappropriately modified.

Please note that:

1. RippleDown User Account data (including hashed passwords) reside in unencrypted form on the server.
2. There is a default Admin account configuration with every new installation. This account should be removed after installation.
3. User names are not case sensitive.
4. Users are required to change passwords after initial logon.
5. RippleDown users are automatically logged out after a configurable period of inactivity. This period should be set to the minimum acceptable value.
6. Depending on the data chosen to send to the RippleDown instance and the manner in which projects are configured, users may come in contact with PHI/PII when using RippleDown.

It is the administrator's responsibility to ensure that the following recommendations are implemented:

1. Robust manual security policies are used when configuring RippleDown User Accounts (for example, complex passwords, regular password changes).

2. The RippleDown user group facility should be used to allocate to users only those specific permissions that they require, that is, the principles of "segregation of duties" and "least permission" should be adopted.
3. Robust operating system-level security policies should be implemented on machines that host RippleDown applications (for example, complex passwords and regular password changes).
4. Physical encryption of hard drives should be used if the drive is removed.
5. Antivirus and antimalware software should be installed on the RippleDown server and client platforms. Additional CPU and memory resources should be provided, if necessary, in order to prevent any degradation in performance caused by the execution of this software.
6. Appropriate Intrusion Protection System and Intrusion Defence Systems should be deployed.
7. Network firewalls should be used to restrict access to the RippleDown server platform to only those ports that are required for communication with the Clinical Data Source and RippleDown clients.
8. All software and services not required for the operation of RippleDown should be removed from the server platform.
9. Minimum periods for Operating System logon timeouts after user inactivity should be set.
10. Administrator access to the Operating System of both the RippleDown server and clients should be restricted to only the authorised admin users.
11. User access to the folder where databases to be backed up are located should be restricted to only authorised administrator users.
12. Backup and restoration of RippleDown databases should be performed in a secure manner by disabling access to physical ports that may allow connections from removable devices. If removable media is required for functional purposes, encrypted removable media should be used.
13. All backups should be encrypted using a secure encryption algorithm.
14. The RippleDown server should be installed on a platform where full disk encryption has been enabled.

> Once encryption is enabled, it is recommended that you validate the performance of RippleDown is not adversely impacted.

15. Physical security measures for both the server and client platforms should be implemented to prevent unauthorised access to these platforms.
16. Robust security policies should be adopted for the repair, removal or destruction of any hardware or network assets associated with the RippleDown application.
17. The data sent from the Clinical Data Source to the RippleDown server contains PII/PHI information hence robust security policies should be adopted to ensure the security of the network and communications between these applications.
18. The data sent from the Clinical Data Source to the RippleDown server should be configured on a per-project basis so that only the minimum amount of data, in particular PII/PHI data, is sent that is necessary for the deployment of that project.
19. The RippleDown log files contain an audit trail of security-related events, including user activity. The system-level access control for the log files should be set such that only an authorised user has access to these log files.

Remote Application Management

Whilst not mandatory, it is recommended that you provide remote access to PKS for the purpose of ease and efficiency in the provision of support. PKS will align with the customer's preferred approach, which may be:

1. Initiated and managed by the customer as required via a technology such as ZOOM,
2. Requested by the customer and initiated and managed by PKS as required via technology such as ZOOM
3. On-demand software-based VPN and remote desktop access, established and managed by the customer but initiated as necessary by PKS or
4. Permanent network-based VPN access and remote desktop access, established and managed by the customer but initiated as necessary by PKS

Remote Access Protocols
At all times, remote access is the responsibility of the customer. The customer should ensure that the connection is secure, the access rights of PKS support staff are appropriate and the data visible to PKS is not sensitive in nature.

Log File Transmission

For security reasons, a RippleDown installation does not, by default, transmit log files from the site.

However, for the purposes of proactive and preventative maintenance, a configuration where daily automatic transmission of logfiles to the PKS support site (or other destination) may be arranged.

If transmission of log files is enabled, the use of a mail server that supports Transport Layer Security is recommended. Please consult the RippleDown Installation Guide for details.

Patient Sensitive Information
At all times, the content of the log files is the responsibility of the customer. Whilst every care has been taken by PKS to ensure sensitive patient data is not reflected in the log files, the risk of this occurring cannot be completely eliminated.

Decimal separator

RippleDown installations by default use (.) as the decimal separator. An installation may be configured to use comma (,) as the decimal separator to allow for customization for users of RippleDown in countries where this convention is implemented.

Please note: The decimal separator must be set at first installation and should not be changed once a knowledge base is created.

The decimal separator is set through the vmoptions for the server and client settings.

Administrator Settings:
• Access the server hosting the RippleDown installation folder.
• Stop the RippleDown service.
• Access the file named 'Client.vmoptions' and amend by adding the line '-Dnumeric.locale=XX' on a new line at the end of the file. XX being the specific local setting, please contact PKS support for specific location settings. Save the new setting.
• Access the file named 'Server.vmoptions' and amend by adding the line '-Dnumeric.locale=XX' on a new line at the end of the file. Save the new setting.
• Start the RippleDown service

Once applied, all evaluations will be made using comma as the decimal separator. This includes syntax as well as attribute reference values.

Please note while using comma as the decimal separator the dot point becomes the indicator for thousands. For examples of values using a comma vs a dot as the separator see table below:

Number Description	Dot as decimal Separator	Comma as decimal separator
1 thousand and 1 quarter	1000.25	1.000,25
3 thousand 4 hundred and 55	3455	3.455
1 half	0.5	0,5

A typical deployment schedule

Step	Activity	Description
1	Install RippleDown	Install RippleDown server and client components.
2	Install Clinical Data Source interface	Install and test the interface from the Clinical Data Source to RippleDown
3	Configure interface	Configure the Clinical Data Source interface to: Send to RippleDown the required data items and historical information required for each Expert or Auditor domain. Receive from RippleDown the interpretive generated reports, autovalidation settings and other workflow actions.

Step	Activity	Description
4	System Test	Perform system testing as per the Technical Service Bulletin (supplied separately).
5	Live “in the dark” testing	Monitor the Production RippleDown configuration for an appropriate amount of time, with reports being sent to the clinical data source by RippleDown, but not yet released to referrers
6	Training	Train users in the operation of RippleDown client applications.
7	Deploy Knowledge Bases	Build and deploy the initial set of Knowledge Bases. Reports may now be released to referrers
8	Validate reports	Commence the process of continuous report validation and periodic update of the Knowledge Bases.
9	Increase autovalidation	Commence the process of increasing autovalidation settings within RippleDown to manage the validation load.