OL Connect Backup licenses

Contents

Introduction

This document explains the basics around application downtime, available options and when it makes sense to use OL's backup license. In addition, it covers some general disaster recovery aspects to help you understand what is and what is not covered by a backup license.

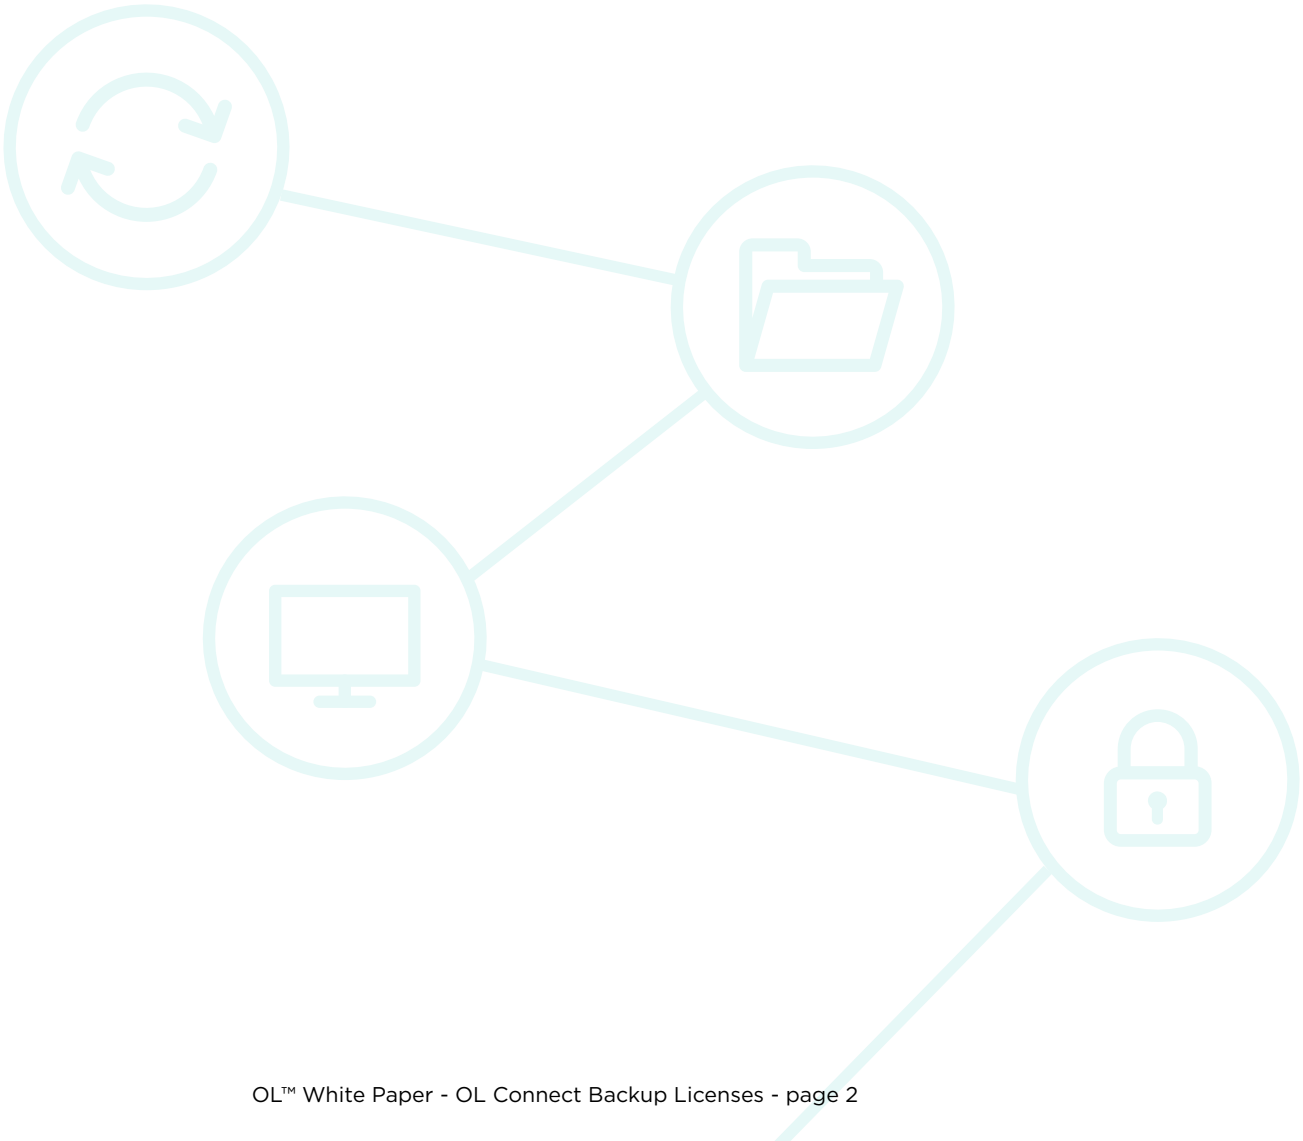

What you need to know about application downtime

Downtime is an important aspect of critical business processes that doesn't often top the budget agenda. In the era of customer experience first, your processes, especially the ones around customer communications, are expected to be available around the clock. High availability can actually be a strategic differentiator that helps you stand out from your competitors. It is a promise you make to your customers that you will deliver, regardless of unforeseen events.

First, it's important to understand that downtime is unavoidable, so focus your resources on minimizing its duration. Many things can cause downtime like hardware failure or environment issues. But, it may come as a surprise how often the culprit is configuration and other human errors. They account for over one-third of incidents[1] (Figure 1). Taking also into account the increasing complexity of technological environments (multiple VMs, servers, etc.). coupled with frequent upgrades and updates, the human factor is not to be taken lightly.

Figure 1 - Percentage of incidents by root cause

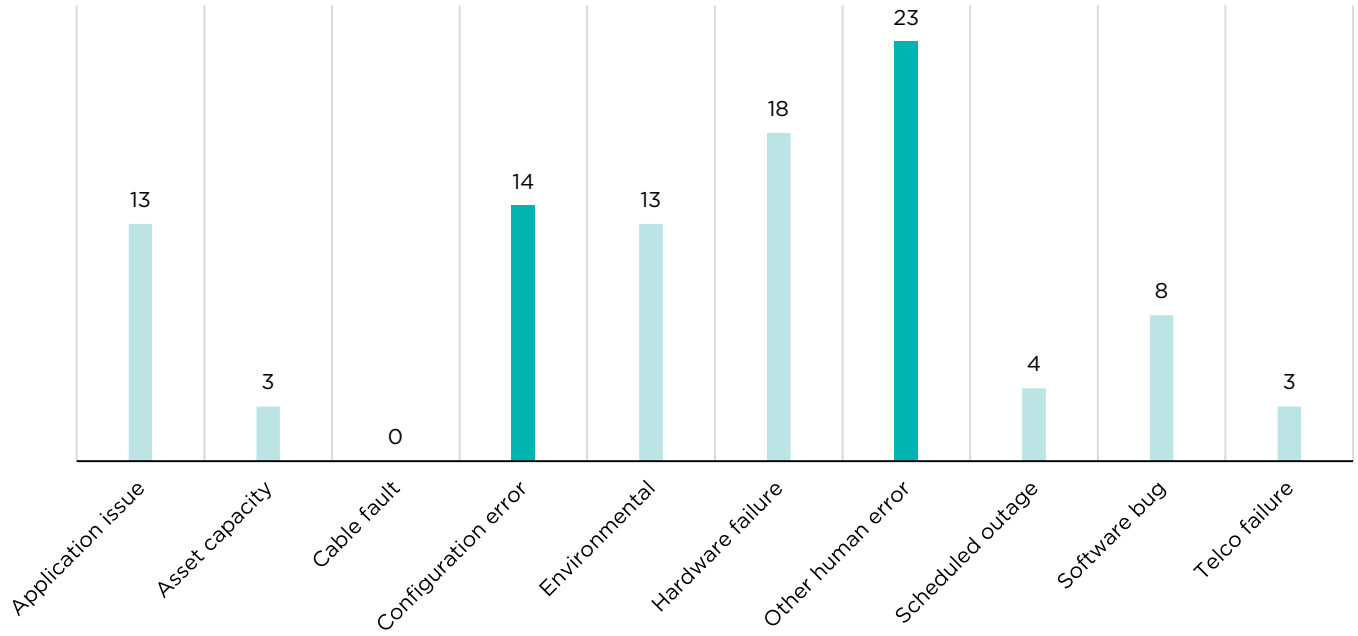

[1] Network Barometer Report 2016, Dimension Data

Regardless of the cause, downtimes have the same disrupting effects:

- Lost revenue - Even a small company can lose on average $8,580 per hour of downtime. This average can easily surpass half a million dollars for a large enterprise [2] (Figure 2).
- Quality of service, customer experience and brand damage
- Possible data loss
- Loss of valuable staff time

The longer the downtime, the higher the cost. **For business critical applications, a study showed an average of 1.33 hours of downtime.**[3] To put this into perspective, a single hour of downtime for a business critical application can be more expensive than the simplest backup plan you can implement.

This is not to say that you should develop a backup plan for every single business process; this is not a viable option for the majority of companies on a tight budget. It is important, though, that you assess your position and the vulnerability of the processes you use or for which you plan to use OL Connect.[4] The backup policy for any company depends on many things, but usually it comes down to the time and the associated cost to the business.

Figure 2 - Average cost per hour of downtime

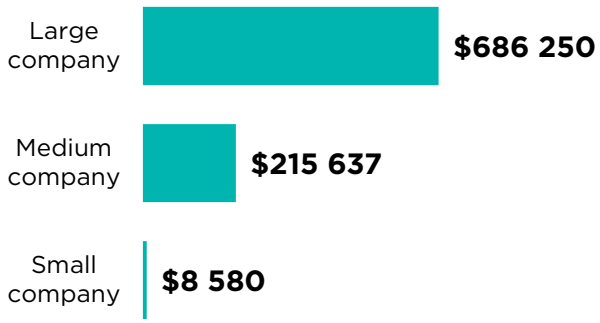

Here are some things to consider before choosing what's best for you:

- What is the impact of OL Connect not being available?
  - If the processes being performed by the production server are low risk, then it could be minimal. **But with half the workload being business critical**,[5] even an hour down, might not be acceptable. In our experience processes are automated with OL Connect because they need to be robust, accurate and reliable. **Such processes are business critical and the cost of downtime can be huge.**
- What are the chances of your server failing?
  - Are you using server hardware or virtualization technology?
  - What if an OS update or installation of third party software leads to unforeseen, even catastrophic, effects?
  - What if a user or administrator accidentally messes up the system beyond an easy repair?
- How long will it take to recover the server?
  - Downtime costs generally put pressure on expected recovery times.
  - IT departments are always busy and generally prefer to be prepared for server failures.
  - Recovery from a live backup server is faster than from an offline backup or even a clone.

[2] Aberdeen, 2016
[3] *Veeam Data Center Availability Report 2014*
[4] Throughout the document we use the term "OL Connect" to mean either PlanetPress Connect or PReS Connect.
[5] *Veeam Data Center Availability Report 2014*

What are my options?

Reinstall, reactivate, and rebuild

When you buy a production license, you also buy OL Care, which allows you to get support and even transfer licenses from one server to another in the event of a failure.

Your production license can be deployed on a different server either by restoring a backup to a different server or by doing a fresh install and configuring the software from scratch. This can take a significant amount of time and will require that you have the resources and skills to rebuild the system.

To continue operating in such a situation, the new server will need to be **activated** with the same serial number. A temporary seven-day activation code can be created online using our web access manager. A permanent license must be obtained by contacting our Customer Care department, and you will be asked to sign a License Transfer Agreement.

Create a Virtual Machine

Your production system could be cloned using Virtual Machine technology. That machine must be offline[6] and can then be enabled in the event that the original server fails.

If a clone snapshot is taken regularly—which requires a backup routine to be implemented—then you can recover much faster because there is no need to reinstall, reactivate or rebuild.

A Virtual Machine is not a high availability backup server, so the backup routine would need to include checks to ensure that any failed jobs are restarted (assuming the clone is installed on a server that is identical to the host).

Run two servers – invest in a backup license

A backup license allows you to install and activate a second license on a backup machine.

The backup server can be activated alongside the production server because it will have its own serial number.

Having a live backup server allows you to create custom workflows that could automatically copy key data and resources to ensure that your backup server is ready to go in the shortest possible time.

So if you want to minimize your backup policy overheads and recovery times, then you need a backup license. Backup licenses are available for PlanetPress Connect and PReS Connect.

A backup license allows you to install OL Connect on an additional server, so it can be available immediately—or at least very quickly—should the production server fail.

When configured with the same options, a backup license is identical to a production license. Backup licenses are only available to those who buy, or have previously bought, a production license.

[6] The clone cannot be online with the original production server because they will use the same serial number, which is not allowed. If the virtual machine is moved to a different host, then it might be allowed.

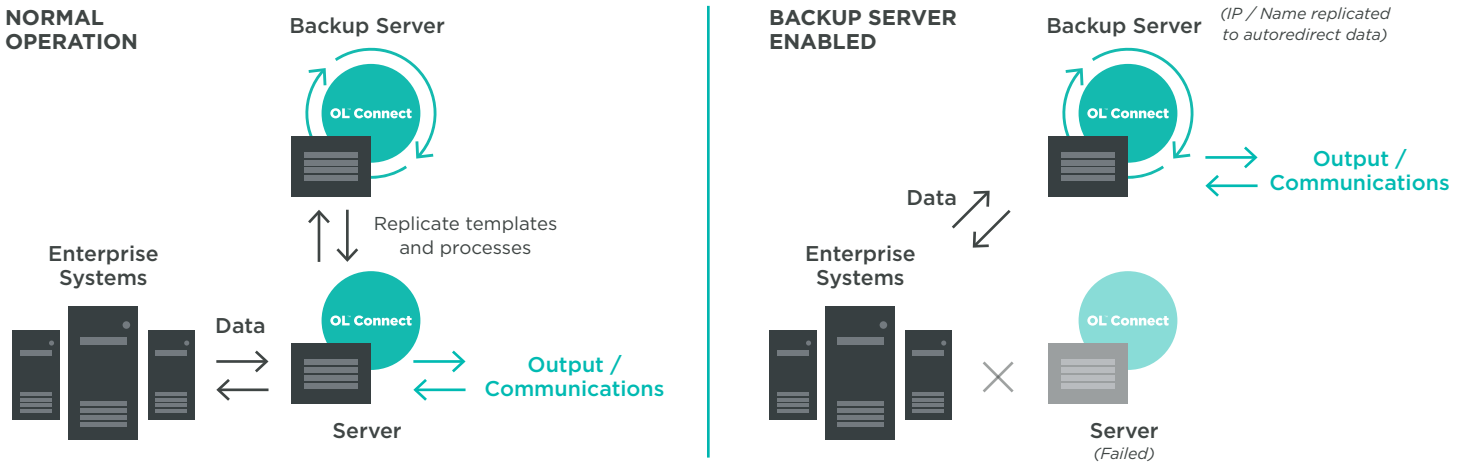

Best practices

Use VM images or snapshots

As mentioned earlier, using virtual machines might be a viable option for you. Here are a few tips to consider:

- Keep an up-to-date copy of the virtual machine image.
- If you have a problem, revert to the copy.
- Make sure the host PC doesn't change or is reconfigured. Otherwise, the activation may fail.
- Expect some data loss and plan for it.

Use a high availability VM platform

If high availability is required, some virtual machine solutions include functionality for moving to a different hardware without downtime. For instance, VMWare vSphere offers vMotion and Fault Tolerance, which are viable solutions.

These solutions protect against hardware failure. Keeping an up-to-date copy or snapshot of a VM also allows recovery from operational and application errors.

Handling data loss and resuming operations

In general, a mid-process failure in an OL Connect Workflow cannot continue from where it left off. So any running jobs will have to be restarted after recovery. It can be important to validate completed jobs in order to minimize missing or duplicated output.

The easiest way to achieve this is to resubmit the job from where it originated. This is similar to what you do when a printer has a problem: you go back to your work station and print again.

Resubmitting a job may not be possible for different reasons: the job is no longer available, some processing steps take too long, or some of the steps have side effects that should not be duplicated. The general approach to deal with this is to design restart points into the processes:

- Copy incoming data to a backup folder as a first process step.
- Split a process into multiple processes that can be restarted independently, if needed. Each partial process can keep a copy of every incoming job.

Preventing duplicated output

Special attention must be given to the point where output is created. When resuming operations, the output for a job that was in the middle of being sent to the printer or being emailed as a batch mailing will be duplicated if you resubmit the job.

Print

In the case of batch print jobs, this can often be handled by the printer operator, who can check which parts of the output were already done. This could result in a lot of paper being discarded or perhaps the printer interface allowing only part of the resubmitted job to be printed.

Email

In the case of email, there is no physical output that can be checked for duplicates. When working with an email service provider (ESP), make sure no duplicate emails are going out:

- The ESP may offer functionality that lets them hold on to your emails until the job is done, or
- The logging/accounting of the ESP may provide the information to establish what emails went out.

Without an ESP, you can send a blind copy (BCC) of every email to an internal server so that the server can be checked for the last email that went out.

In either case, the OL Connect Workflow processes will have to allow sending out only part of an email job.

Test your recovery scenario

Whatever choices have been made for you may run into problems if you try to switch production, or you may not be able to switch within the required time.

How often should this be tested? That depends on the amount of changes that are typically done to the server configuration, document templates, etc. If a system configuration is fairly static, with hardly any changes to processes and templates, then once a year may be sufficient. If you are updating templates and changing processes and other configurations on a weekly or monthly basis, then you should test your backup server more often, perhaps every month, week or day.

If you find that testing your backup scenario is too disruptive to your business processes, this may be a sign that the availability of your backup server is too low. For instance, you may need to switch from cold standby to hot standby.

Keep copies of configurations, resources and the runtime environment

Having a server standing by is not much use if that server is not properly configured. You need to make sure that you have the production versions of your configuration files, document templates, etc., available. If these reside only on your production server, you may lose them together with your server.

If switching to the backup server has to be (nearly) instantaneous, the backup server has to be kept synchronized with the production server. Depending on your situation, this can be done in one of two ways:

- Manually – Every change to the production server is also applied to the backup server. This requires discipline, but can be sufficient for environments with infrequent changes.
- Automatically – For instance, an OL Connect Workflow process can automatically deploy modified resources on both servers.

Most OL Connect applications have some kind of runtime environment. This can be as basic as some hot folders, but it can also include runtime status files, and additional resource files. Keeping copies of these can be as simple as periodically creating a zip file (manual or automatic). Alternatively, files and folders can be stored on redundant networked storage, which allows the backup server to use the same environment as the production server.

Q&A

I want to set up two servers for automatic failover. Do I need a backup license?

Yes.

Does a backup license include automatic failover?

No, it doesn't. You simply get the capability to have a second server active in the same network segment. In general, it is possible to configure these two servers for automatic failover by configuring this into the automation processes of OL Connect Workflow. In addition, the server infrastructure should be set up appropriately, so jobs either get resubmitted (to the backup server) in case of failure, or they are saved in shared storage so the backup server can pick them up from there. Depending on the application, additional infrastructure such as load balancers for http requests may be required.

PReS Connect offers clustering. Does that help with disaster recovery and failover?

No, it doesn't. The clustering functionality for PReS Connect is for performance clustering. The Server Extensions in the cluster help the main server to carry out its tasks, but they can't take over from the main server.

The kind of clustering that is used for automatic failover is called high availability clustering, or HA clustering. This is different from the performance clustering that PReS Connect offers. OL Connect's performance clustering is more like high performance clustering, or HP clustering, such as can be found in super computing.

objectiflune.com