MyQ X

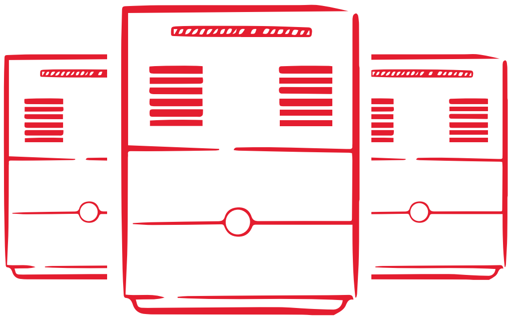

High Availability

Whitepaper

November 2022

Table of Contents

1. Introduction

Access to IT, data, and services is essential in an organization's day-to-day activities. Organizations have often become reliant on IT systems, for example, the availability of Print Management services. Downtime of IT systems may lead to adverse business impacts, and as a result, there is a requirement for IT systems to be more dependable.

This whitepaper provides an overview of MyQ's approach to achieving Higher Availability and IT system continuity in terms of Print and the components related to Print Management via MyQ X. It is essential to note that different organizations may attach importance to various aspects of achieving Higher Availability and IT system continuity, which is dependent on the needs and requirements of the respective organization. Additionally, we understand and embrace the fact that IT environments and architectures may differ between organizations. Therefore, MyQ adopted a universal approach in dealing with the subject of High Availability and IT system continuity.

When it comes to High Availability, an organization needs to consider several factors before choosing the most suitable HA implementation. The first question that needs answering is **what needs to be protected to ensure IT continuity** and which level of High Availability meets the organization's requirements. Regarding what needs protection in this white paper, we will focus primarily on components required to ensure a higher level of continuity in delivering Print and Print management via MyQ X.

In an ideal world, IT professionals would like to have their IT systems protected with the highest level of Availability (99.999 percent); however, this is challenging and costly to achieve. Even at the highest level of High Availability, IT continuity is not guaranteed 100 percent of the time.

You may face two types of downtime in your organization: planned and unplanned. Planned downtime may result from maintenance, applying patches, updates to software, system maintenance, etc. This type of downtime is usually unavoidable and required for optimal system performance and stability. Unplanned downtime is due to an unforeseen event, for example, a hardware or software failure.

Consider that the higher the level of Availability, the higher the cost implications will be. As an organization, you need to justify the decision from a financial point of view.

Additionally, most organizations have already invested in High Availability implementation to protect their existing IT systems, investment in financial cost, time, and training of IT staff to maintain and manage the implementation. In these cases, it would not make logical sense to convince the respective organization to implement some propriety solution for High Availability that they are not familiar or comfortable with. **MyQ's approach accommodates these high availability environments and provides organizations with unique tools and features to complement existing environments.**

2. High Availability

2.1 What is High Availability?

In simple terms, **High Availability is the ability of a service to continue operating despite failures within its environment.** You can achieve HA by designing a system with no single point of failure. In a high availability system, workloads are distributed across a cluster of server nodes. If one server node fails, the workloads running on it automatically move to other servers. It helps ensure that critical applications keep running even when problems occur in other parts of the system.

The main goal of High Availability is to **minimize the impact that downtime/outages have on an organization's business processes.**

A common misperception of High Availability is that an IT system will be available 24/7. Even at the highest level of HA, an organization should still prepare for a small percentage of planned or unplanned downtime, which an SLA usually defines with "Five Nines" being the highest level of Availability.

2.2 How much High Availability do you need?

If you were to ask people how much their IT systems should be protected against downtime, the obvious reply would be 24/7, always available; however, this is incredibly challenging and costly to achieve.

Before any organization can consider a strategy for High Availability, a proper **business impact analysis** is needed to identify critical business processes and the risk related to planned and unplanned downtime of the interconnected IT systems.

Although some organizations may require the highest availability level, many do not. Additionally, **not all IT services are business critical; therefore, it may be more acceptable and economical for some organizations to have a High Availability SLA of 99% vs. 99.999%.**

A simple example could be a Supermarket that requires its online payment system to be available at the highest level of Availability. The impact of not having this service available can result in loss to the business because customers will not be able to make any purchases if they do not have cash on them. The same supermarket will, however, be less impacted if the print system in the back office were to go offline for a few minutes due to planned or unplanned downtime.

2.3 Determine High Availability requirements with a BIA (Business Impact Analysis)

Before deciding on a High Availability strategy, an organization should carry out a BIA (Business Impact analysis) to determine the severity of impact that IT-related outages/downtime have on critical business processes and to identify requirements needed to ensure continuity of operations.

Business impact analysis may differ from organization to organization but usually covers the following:

- Identifies critical business processes
- Calculates a measurable risk of loss due to IT outages/downtime
- Considers essential business functions, people, and business dependencies
- It is based on data gathered via BIA interviews with employees

Ultimately the BIA will allow the organization to see how a business would be affected if you took business processes away during IT systems outages/downtime. It helps the organization determine which business processes are the most critical for continued operation and assists in creating a recovery plan.

2.4 Key Recovery Objectives

When considering High Availability, there are two key parameters that define how long your organization can afford to be offline and how much data loss it can tolerate. These parameters are the **Recovery Time Objective (RTO)** and **Recovery Point Objective (RPO)**

- **Recovery Time Objective (RTO)** is defined as the maximum length of time it takes for an IT system, set of applications etc. to recover from downtime (planned, unplanned or a disaster) and resume standard business operations. RTO timelines are decided amid Business Impact Analysis (BIA) during business and IT continuity planning.

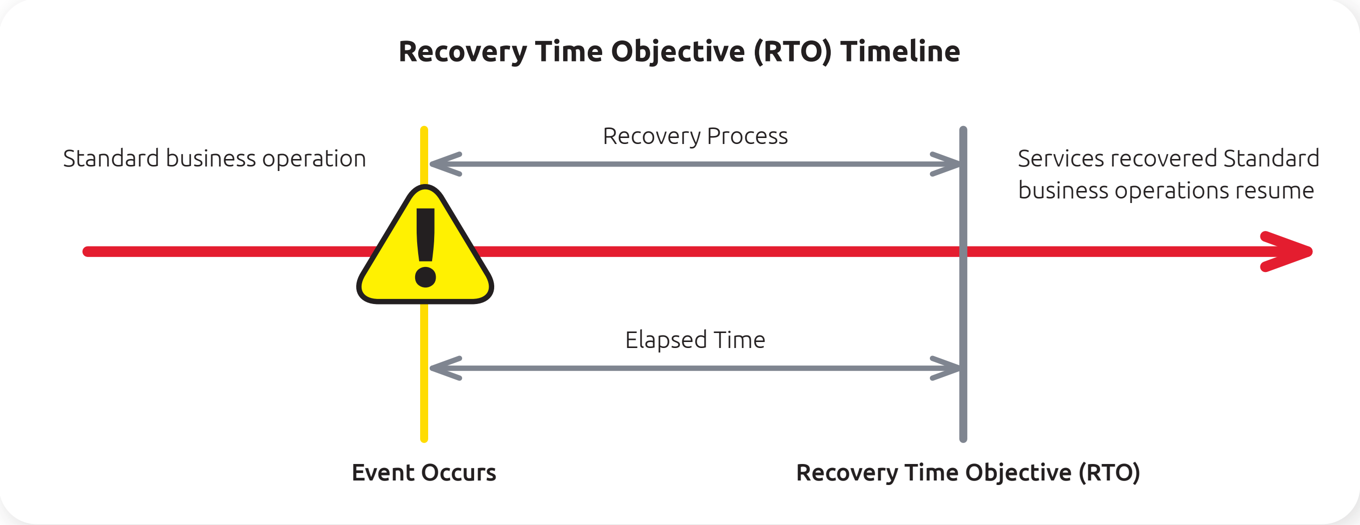
Recovery Time Objective (RTO) Timeline
Standard business operation
Recovery Process
Services recovered Standard business operations resume
Elapsed Time
Event Occurs
Recovery Time Objective (RTO)

Figure 2.4.1: Illustration RTO timeline

- **Recovery Point Objective (RPO)** is a measure of how often you back up. Can you afford to lose a certain number of minutes, hours, or days of data updates if disaster strikes between backups? RPO indicates how recent the restored data will be. For example, if you experience a failure now and your last full data backup was 24 hours ago, the RPO is 24 hours.

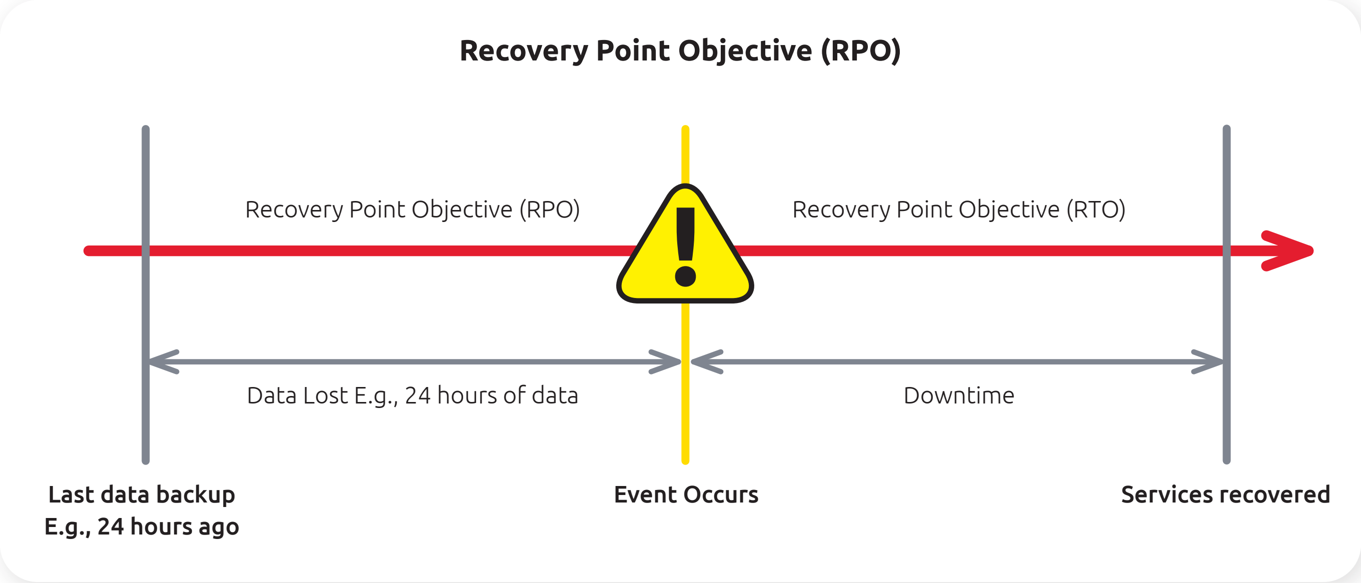
Recovery Point Objective (RPO)
Recovery Point Objective (RPO)
Recovery Point Objective (RTO)
Data Lost E.g., 24 hours of data
Downtime
Last data backup
E.g., 24 hours ago
Event Occurs
Services recovered

Figure 2.4.2: Recovery Point Objective (RPO)

Essentially, RPO has to do with the frequency of backups, while RTO refers to recovery time.

2.5 How to measure High Availability

High Availability is measured in the percentage of time that a service is available to users, often referenced by the number of nine's in the digits. "Five Nines" is used to describe an IT system's continuity with 99.999 uptime. In other words, the IT system or service is only unavailable for 5.39 minutes throughout the year for planned or unplanned downtime.

Number of 9s	Percentage of Uptime	Amount of Downtime (Year)
Two 9s	99%	3 Days 15 Hours
Three 9s	99.9%	8 Hours 45 Minutes
Four 9s	99.99%	52 Minutes 33 Seconds
Five 9s	99.999%	5 Minutes 15 Seconds

Achieving five-nines of High Availability over some time is incredibly challenging. It is expensive due to the running costs of the physical hardware infrastructure and software components, and additional components add to complexity and risk. For many services or networks, three or four nines would be more effective and justified regarding the resources and cost involved.

3. Achieving High Availability with Application Clustering

3.1 What is Application clustering?

Application clustering (also known as software clustering) uses software to configure multiple independent computer systems, referred to as nodes, into a cluster. The cluster of nodes then works together as a unified computing resource. A benefit of Application clustering over its hardware-based counterpart is that nodes can easily be added or removed from the application cluster as and when required. Application clustering is, therefore, more scalable. Furthermore, because it does not require specialized hardware, an application cluster tends to be more economical and less complex to configure. For these reasons, application clustering is the commonly preferred method.

Benefits of deploying MyQ X with application clustering

Increased resource availability – In a cluster of servers (nodes), resources such as compute power, memory etc. are shared. If one of the nodes requires maintenance or downtime, the other nodes can take over the workload.

Increased performance – In a cluster compute resource of multiple nodes are combined to provide greater processing power than a single server can deliver on its own.

Greater scalability – As your userbase and requirements grows you can easily assign more compute resources as required.

Reliability and Failover Support – Application clusters usually have a central system which continuously monitors the health of the nodes within the cluster. When there is a failure of a node within the application cluster the system automatically assigns workload over to the remaining nodes. When combined with a shared storage solution data is safeguarded during failure. There is little to no downtime for the user during the transition of workload from one node to another.

In summary, **by deploying MyQ X within an application failover cluster, you can protect your print and document services from planned and unplanned downtime.** Secure print and document data. Keep users productive as there is no interruption of service during downtime.

3.2 Active-Active vs. Active-Passive configurations

The typical **active-active** cluster consists of at least two nodes, and both actively run a similar service simultaneously. The main objective of an active-active cluster is to achieve some load balancing. Load balancing is distributing the workload across all available nodes evenly. This concept prevents any single node from overloading because multiple nodes can handle the workload. Additionally, load balancing provides an improvement in throughput and response times.

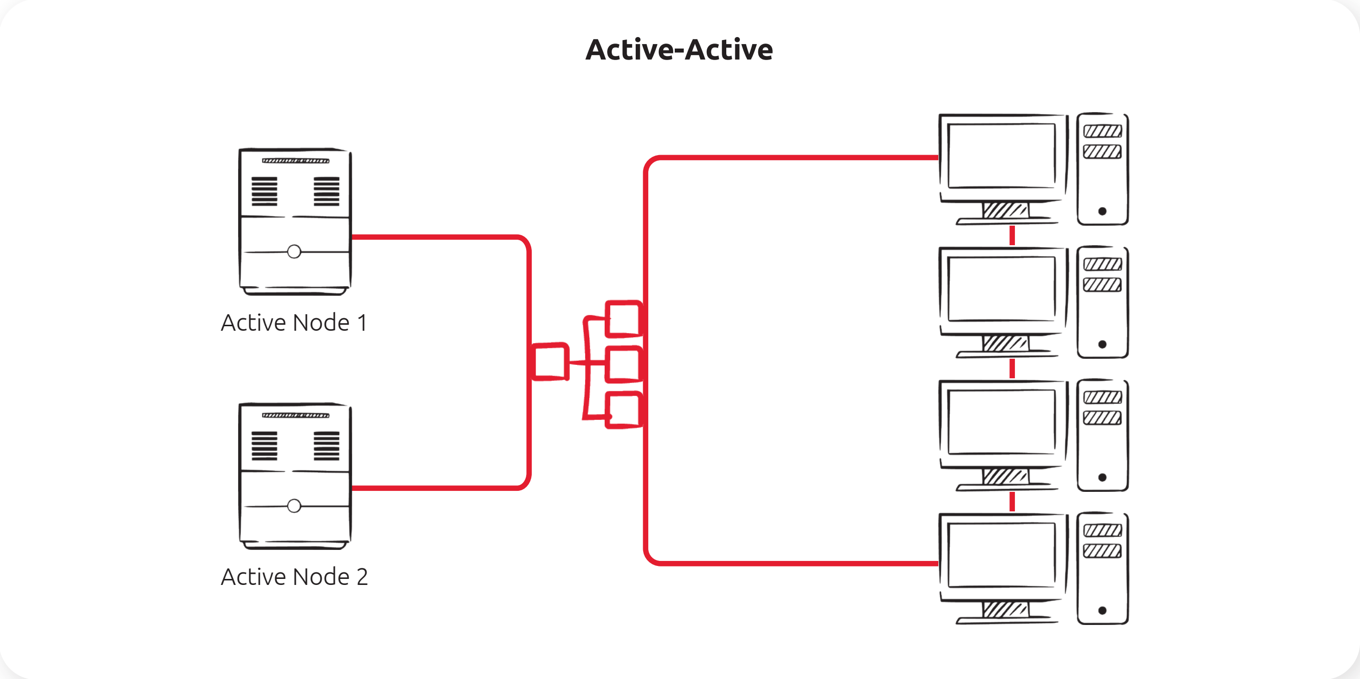

Figure 3.2.1: Example of an Active-Active cluster configuration.

The typical **active-passive** cluster consists of at least two nodes. However, not all nodes are active in an "active-passive" configuration. For example, in a two-node design, the first node will be active, and the second node will be in passive or standby mode.

The passive (failover) server is a backup ready to take over if the active (primary) server disconnects or when the node fails.

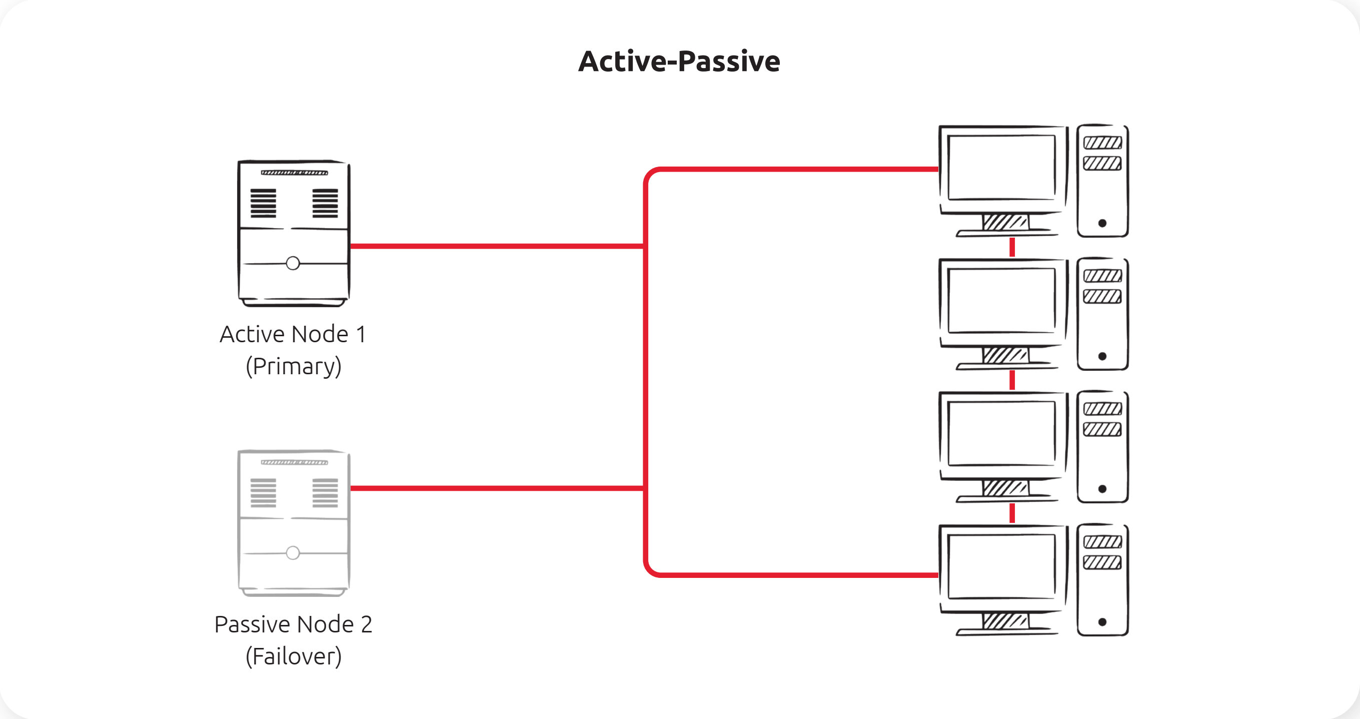

Figure 3.2.2: Example of an Active-Active cluster configuration.

3.3 Overview of Active-Active vs. Active-Passive configurations

Active-Active	Active-Passive
Provides a Higher level of Availability compared to Active-Passive	Provides High Availability
Failover is instant as the already running Active nodes will just take over the load of the failed node.	Failover can take a few minutes (dependent on several factors) as the Passive node and services are started up from standby mode to assume the role of Primary node.
Requirement of Load Balancers, it needs to be considered that independent load balancers can also fail.	Active-Passive implementation does not have load distribution as the passive node is running in Standby mode.
Active-Active implementation is more complex to implement and maintain. Requires high investment in IT staff.	Less complexity vs Active-Active, easier to maintain.
Due to additional components and resources Active-Active Implementation is considerably more expensive vs Active-Passive.	The passive node is running in Standby mode most of the time the configuration is more cost-effective vs Active-Active.
System performance suffers during fail over as the load placed on the remaining nodes increase.	System performance remains the same during fail over.

4. Protecting MyQ X with High Availability

4.1 High Availability Implementation with MyQ X

Taking into consideration what we have discussed up until now, **we will continue to look at different ways of protecting MyQ X with High Availability technologies, Clustering, and MyQ X's unique set of failover features.** By this point, you should have a clear overview of your business requirements relating to the Availability of Print and Print management services.

We know that there are many HA methods and technologies available that your organization and IT may already be familiar with and currently use in your existing infrastructure. These HA methods and technologies may deliver similar functionality to that which we will be covering, such as the Microsoft Failover Cluster and VMware's virtual cluster. In our examples, we will explain the core concept of achieving High Availability with MyQ and your print environment.

If you can achieve the same result with your existing technologies, **we recommend that you use what your IT has been trained in and what you are familiar with.** When using your own methods and technologies we do, however, recommend that you deploy a proof of concept in a controlled environment and do **thorough testing before deploying to a live environment.**

At MyQ, we also considered that there might be smaller organizations out there that also require print services to be available during downtime. These organizations do not necessarily have the budget or trained IT staff to deploy expensive and complex HA methods and technologies. Therefore, **we have developed a set of MyQ X failover features that are free to all MyQ X customers. These features were designed to keep our customers printing during planned or unplanned server downtime.**

4.2 Microsoft Failover Clustering

Microsoft Failover Cluster is a Windows Server (Operating System) based feature. The software solution facilitates the grouping of multiple independent hardware servers (thereafter referred to as nodes) into a cluster. The nodes in the cluster collectively work together to provide advanced capabilities such as resource management, health monitoring and failover coordination.

In terms of storage requirements for high availability implementations, it is recommended that **there should be at least one shared storage location that is accessible to all nodes in the cluster.** The interconnected nodes within the cluster should be configured to access the shared storage via, for example, a high-speed network connection.

The interconnected communication between the multiple nodes combined with shared storage makes this solution highly available. If one of the clustered servers (nodes) fails, the other nodes will take over to provide service (the failover process). The clustered nodes are proactively monitored. **When the software detects that a node is not working correctly, they are automatically restarted or moved to another node.**

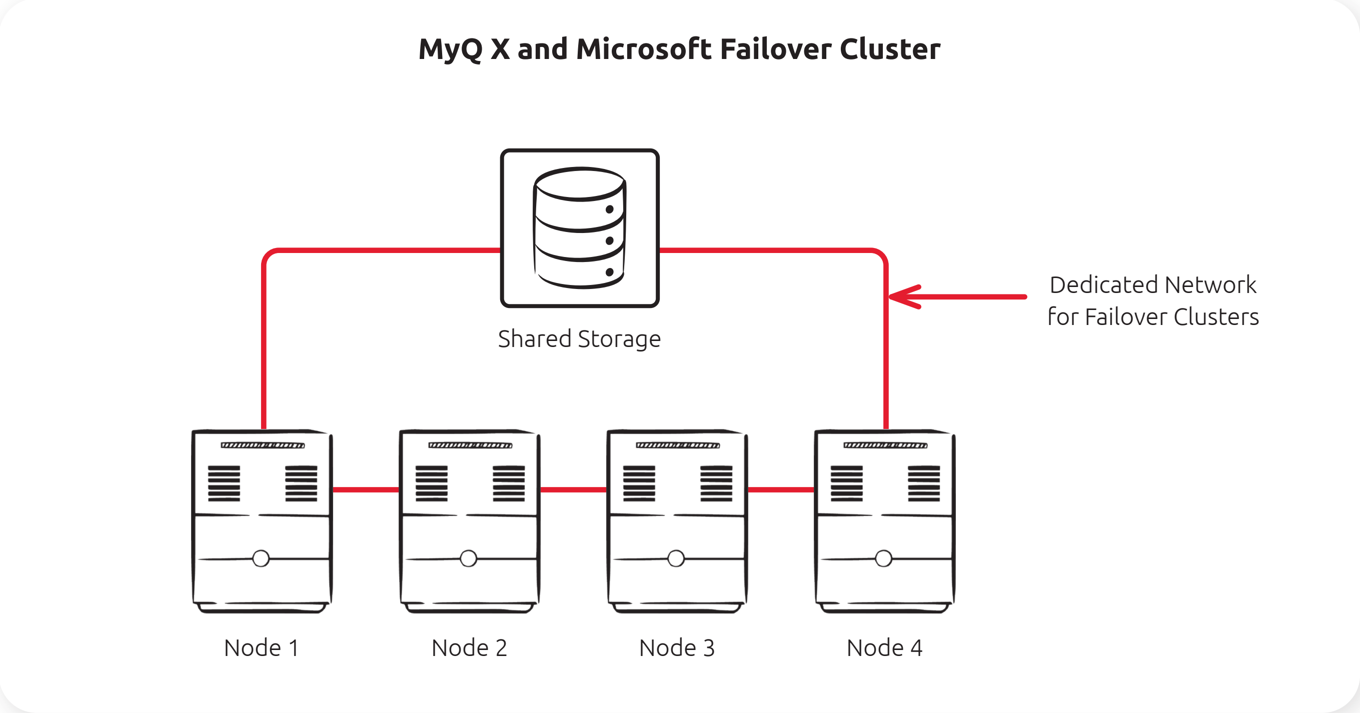
MyQ X and Microsoft Failover Cluster
Shared Storage
Dedicated Network for Failover Clusters
Node 1
Node 2
Node 3
Node 4

Figure 4.2: Illustration of Microsoft Failover Cluster

4.3 MyQ X and Microsoft Failover Clustering

The MyQ X and Microsoft Cluster high-availability concept consists of grouping at least 2 nodes into a failover cluster. The failover cluster will be managed using Microsoft's Windows Server failover cluster functionality and should be configured for shared resource management, health monitoring and failover connection.

The nodes within the failover cluster must be configured to **include a high-speed network attached shared storage that is accessible by all the interconnected nodes** within the failover. MyQ X is then deployed within the cluster in an active-passive configuration with the MyQ X Print server installed on at least 2 nodes. Microsoft's failover cluster software administrates the MyQ services, if the currently active node becomes unavailable, it switches to one of the available passive nodes.

The benefit of deploying MyQ X within the Microsoft failover cluster is that **the system is no longer dependent on only one physical hardware server and is therefore more resilient towards hardware failures.** Read more: MyQ and MS Cluster

Example

In the illustration below **Figure 4.3** the MyQ X Print Server is deployed in an **active-passive** configuration within the Microsoft failover cluster, Node 1 set as active and Node 2 set as passive. Node 1, Node 2, and the shared storage are interconnected over a dedicated high-speed network. In addition, this configuration also requires a separate network that interconnects the user clients and multifunctional printers with Node 1 and Node 2. The separated networks are based on the high availability concept to eliminate single points of failure

Microsoft's failover cluster software administrates the MyQ services, and continuously monitors the health of the nodes. If it detects that the active node (Node 1) becomes unavailable it will try and recover the services, if it is unable to recover the services it automatically switches to the passive node (Node 2).

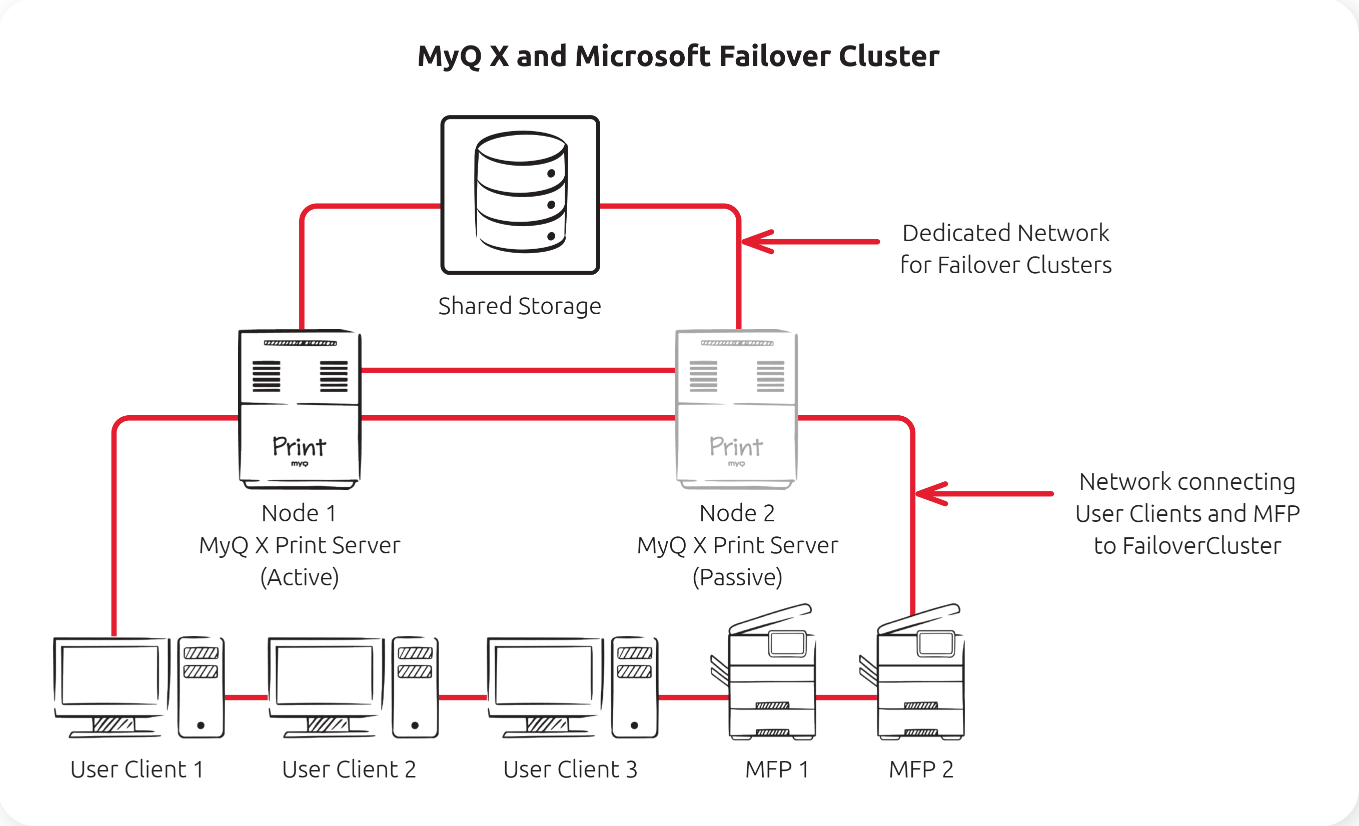

Figure 4.3: MyQ X deployed in Microsoft Failover Cluster

4.4 Virtual Machine Clustering (VMware Cluster)

The primary purpose of **Virtual machine clustering is to run services in a virtualized environment, ensuring high Availability and enhanced server utilization.** Hardware failures within a virtual machine cluster have minimal impact on virtual machines as the VM will transfer to another node.

A Virtual machine cluster typically consists of several components. You will need hypervisor technology, for example, VMware's vSphere ESXi. The vSphere ESXi image is installed on bare-metal hardware just like an OS and acts as a Type 1 hypervisor (referred to as an ESXi host). The ESXi hosts are configured so that their resources are shared.

An essential component of configuring a High Availability virtual machine cluster is to choose an adequate storage solution that is highly resistant to faults with no single points of failure. **MyQ recommends combining your virtual machine cluster with a network attached shared storage solution.** However, different virtual machine clustering software providers deal with cluster storage in various ways and there may be other storage solutions that better fit your organization. The general concept of eliminating single points of failure should remain, the storage solution that you choose should meet your induvial business requirements in terms operational uptime and cost.

ESXi hosts can be managed using a centralized server management platform like VMWare vCenter. vCenter allows for central management of the underlying compute resources of the ESXi hosts and administration of virtual machines, including a host of features such as VM Cloning, High Availability, Fault Tolerance etc. In addition, using vCenter, you can group several ESXi hosts to create a virtual machine cluster. Each virtual machine in a cluster is interconnected via a virtual network.

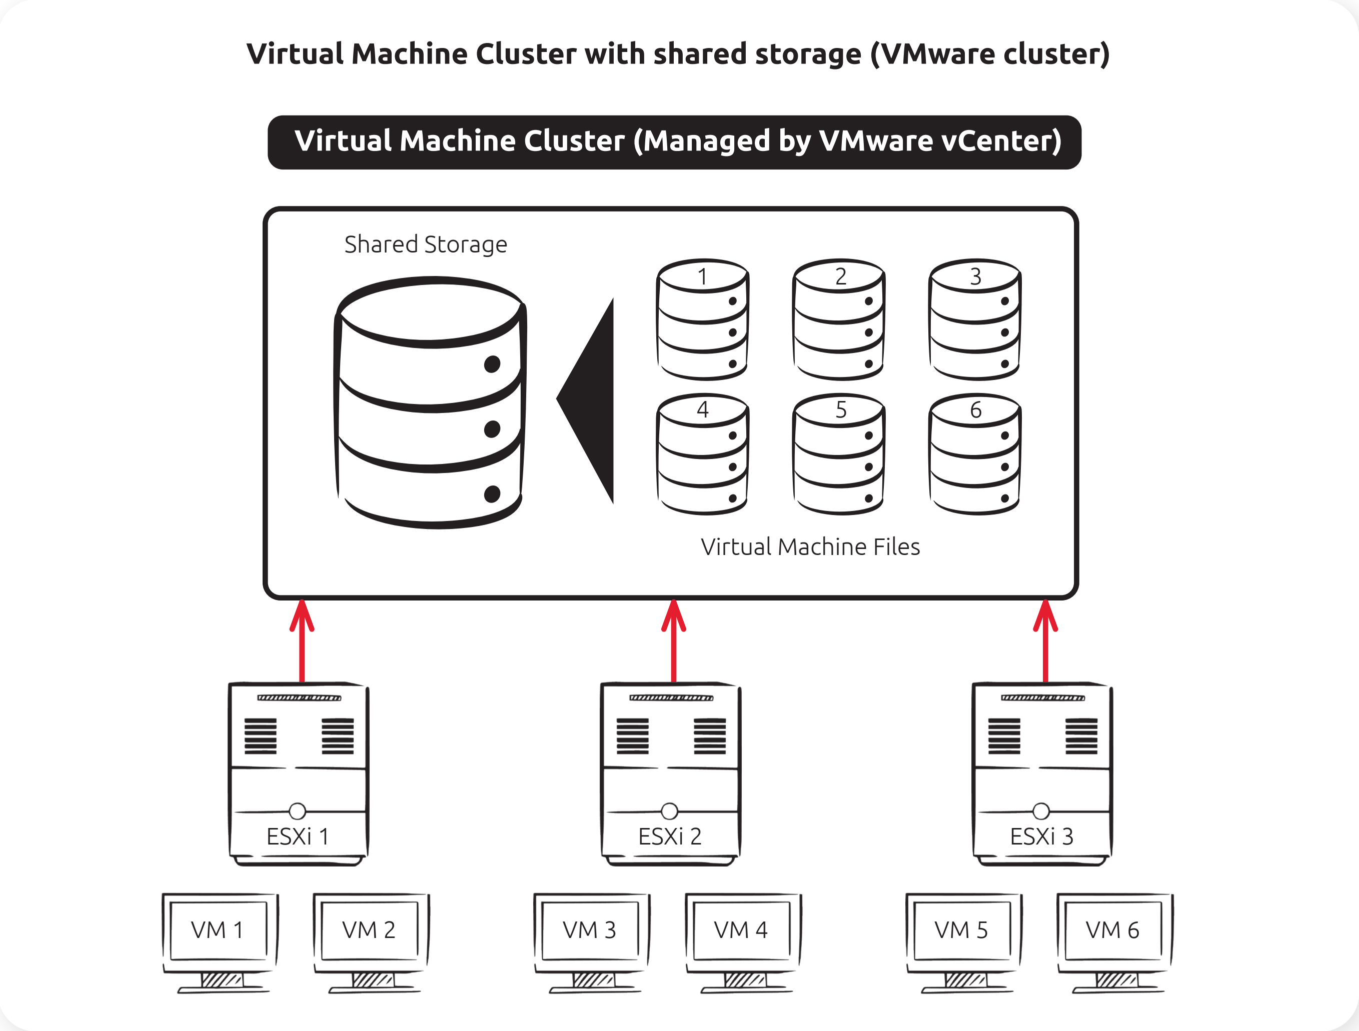
Virtual Machine Cluster with shared storage (VMware cluster)
Virtual Machine Cluster (Managed by VMware vCenter)
Shared Storage
1
2
3
4
5
6
Virtual Machine Files
ESXi 1
ESXi 2
ESXi 3
VM 1
VM 2
VM 3
VM 4
VM 5
VM 6

Figure 4.4: Example of a Virtual machine cluster using VMware technologies.

4.5 MyQ X and Virtual Machine Clustering (VMHA)

The MyQ X virtual machine clustering concept also referred to as VMHA (virtual machine high availability mode) consists of **running the MyQ X system in a virtualized cluster environment in combination with NAS (network attached storage)** datastore.

The NAS (network attached storage) consists of external storage systems which the server nodes (ESXi hosts) will use to store and access virtual machine files remotely. The ESXi hosts will access these systems over a high-speed storage network. The datastores on networked storage can be accessed by multiple server nodes (ESXi hosts) concurrently. The MyQ X print server is installed on a virtual server node (ESXi host), with the workload and resources of the virtual node distributed across the virtual cluster. When utilizing MyQ X's integrated system database (Firebird database) the data will be stored on the virtual instance. With this concept data are always available whenever the MyQ X virtual instance is moved to other nodes within the virtual cluster. In addition to database data, print job data is stored within the jobs folder on the MyQ X virtual instance.

The benefit of deploying MyQ X with VMHA mode in a virtual cluster is that **the system is no longer dependent on only one physical hardware server** and is therefore more resilient towards hardware failures. **VMHA can also operate in a private cloud environment like, e.g., MS Azure**, with the prerequisite there are established VPN (Virtual Private Network) tunnels between the MyQ X server and the organization's local network.

Virtual Machine Cluster failover operation

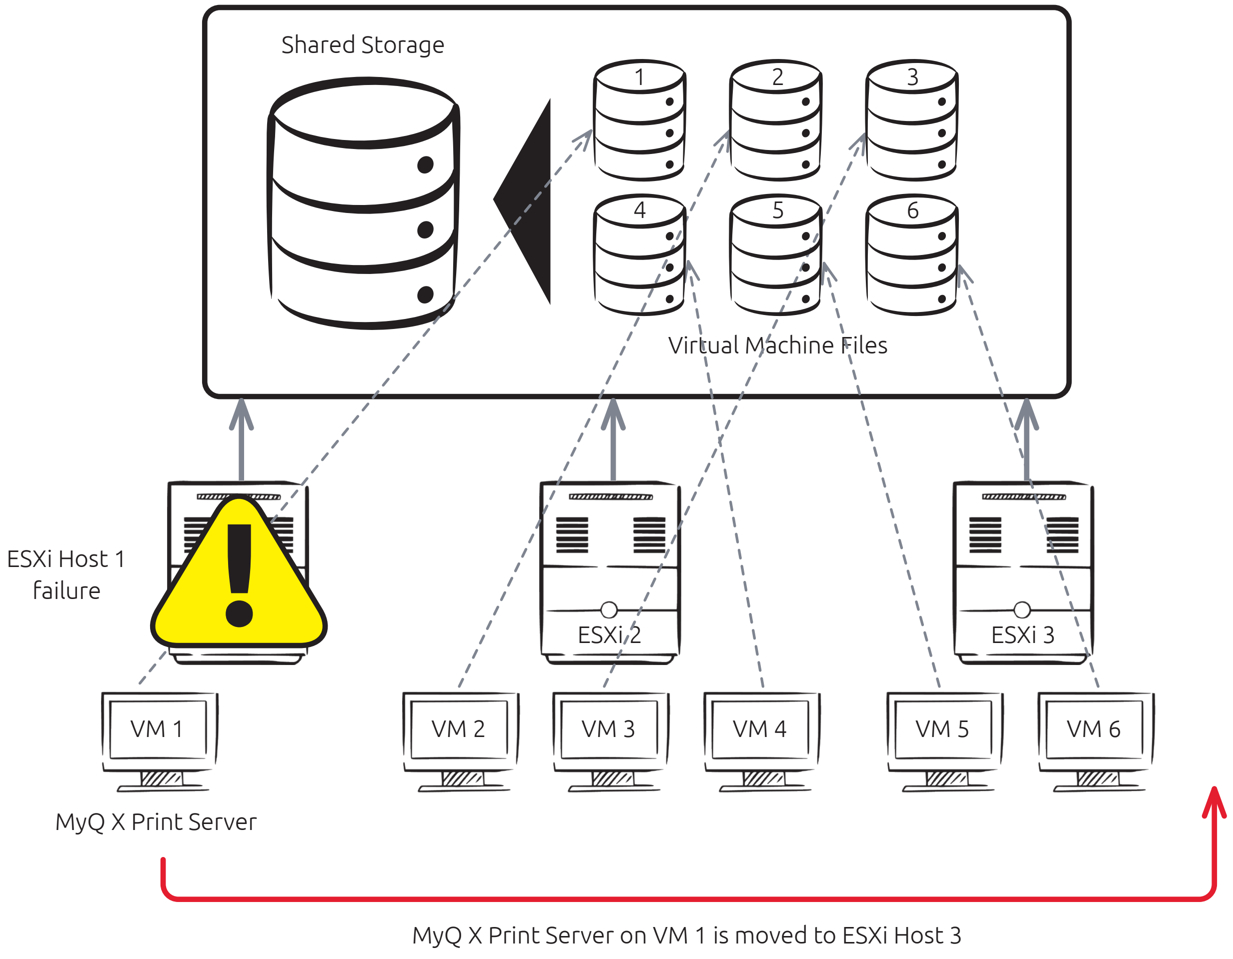

Figure 4.5.1: Virtual machine cluster failover operation

Example

In illustration Figure 4.5 the MyQ X Print Server is deployed on VM 1 which is assigned to ESXi Host 1. The virtual machine image of VM 1 is stored securely in the shared storage location, the shared storage can for example be a NAS (Network attached storage). If ESXi Host 1 fails VM 1 (MyQ X Print Server) will move for example to ESXi Host 3 and continue working uninterrupted. The MyQ X Print Sever data and print job remains secure in the datastore and unaffected by the hardware failure.

5. MyQ X Failover Features

With MyQ X we provide a set of MyQ X failover features that are free to all MyQ X customers. These features are designed to keep organizations printing during planned or unplanned server downtime.

These features can be used to add an additional layer of failover that complements all types of MyQ X implementations, whether it be on-premises, on physical hardware servers or even as an addition to your existing High Availability architectures.

5.1 What is MyQ X Desktop Client?

MyQ X Desktop Client is a complimentary software client to the MyQ X print server that is installed on the MyQ X users' Windows or macOS workstations. MDC provides additional features to users, such as user identification, secure encrypted communication between user and the server, local job parsing, local job accounting, interactive job processing, alternative failover printing methods, monitoring of locally attached printing devices and much more. We will discuss some of these features in the context of High Availability and failover in more detail. Read more: MyQ Desktop Client for Windows, MyQ Desktop Client MacOS

5.2 MyQ X Fallback printing with MyQ Desktop Client

With the MyQ X Desktop Client installed and running on the end user's **Windows** or **macOS** workstation, you can specify a backup printing device via IP/Hostname or choose from a selection of previously used print devices when a connection to the MyQ X Server is lost. When MDC detects a lost connection to the MyQ X print server, it will direct the print job to the specified or user-selected fallback printing device.

Once the connection to the server is re-established, the job accounting will be updated automatically. **The benefit of Fallback printing is that users will still be able to print during planned or unplanned downtime of the server.** Take note that services like scanning and advanced workflows which require an active server connection will not be available during downtime.

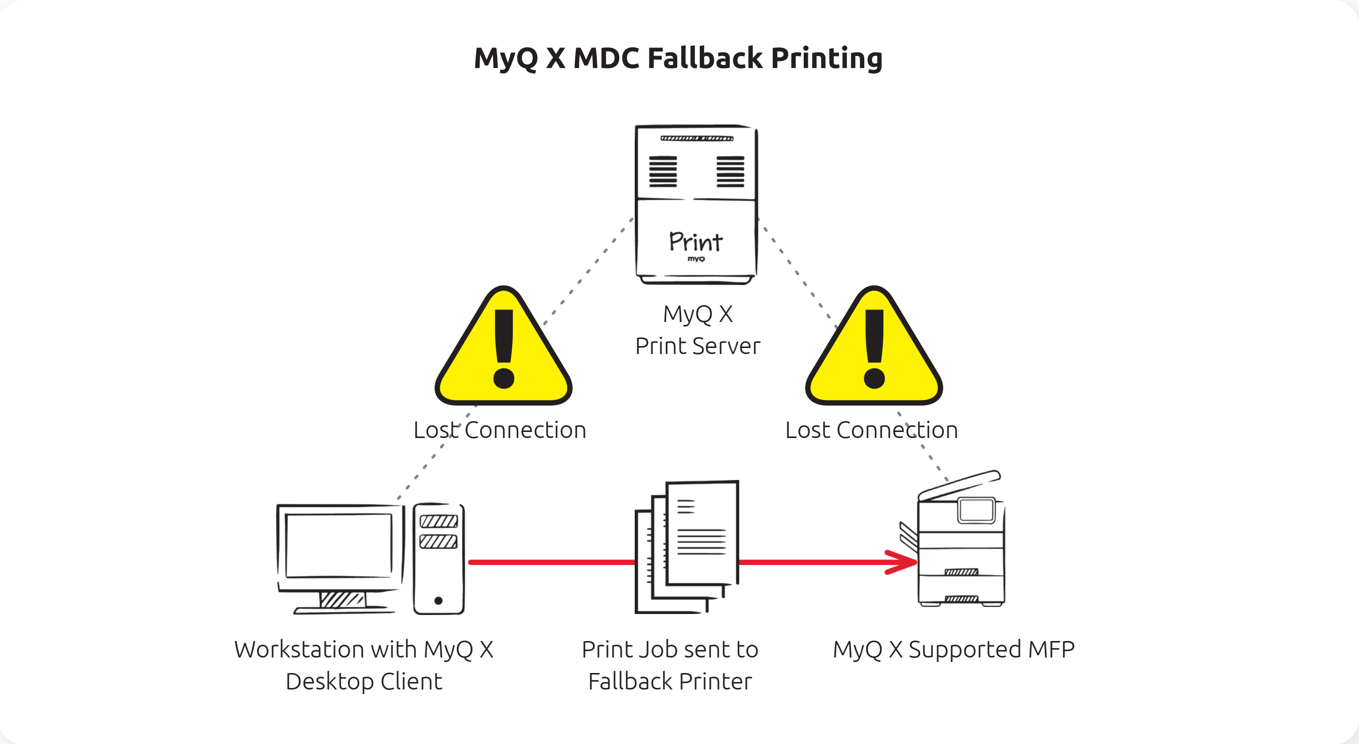
MyQ X MDC Fallback Printing
Print
MyQ X
Print Server
Lost Connection
Lost Connection
Workstation with MyQ X Desktop Client
Print Job sent to Fallback Printer
MyQ X Supported MFP

Figure 5.2: Fallback Printing with MDC when the connection to the MyQ X server is lost

5.3 MyQ X Device Spooling and Offline Login

MyQ X offers availability & resiliency directly via the MyQ X embedded the MFP with our unique **Device Spooling** and **Offline login** functionality. These features are supported directly via MyQ X's intuitive embedded terminal and does not require additional installation of a software client on the end user's workstation.

Offline login is a MyQ X embedded terminal failover feature that **allows users to authenticate on an MFP when a connection to the MyQ X server is lost.** The MyQ X embedded device automatically caches a user's last used login data (e.g., PIN or password). On supported devices, there is also the option for automatic synchronization. The MyQ X system automatically uploads up to 100 user accounts, including their login credentials, to the memory of the selected MFP.

Device spooling is a MyQ X embedded terminal feature that **allows users to continue printing when a connection to the MyQ X server is lost.** The MyQ X embedded terminal supports the functionality natively and does not require the MDC client to be installed on the user's workstation. With device spooling, the user sends the print job directly to the device, which is processed directly by the MyQ X embedded terminal. Once the MyQ X server connection is re-established, the MyQ X embedded terminal will update the respective accounting data.

Device spooling supports several print-release options;

- **Direct printing:** The job prints automatically after the printing device receives it
- **Hold Print:** The job is received by the printer and waits there until the user logs in and prints it. It is impossible to release this job on a printer other than this dedicated printing device. Hold Print is a secured release. It does not share the job with other devices.
- **Pull-print:** The job is spooled by the printing device. When the user logs on to other devices connected to the same subnet, the information about this job is provided and displayed in the list of the available print jobs.
- **Delegated:** It works the same as the device spooling pull-print, except that delegates of the sender can print the job.

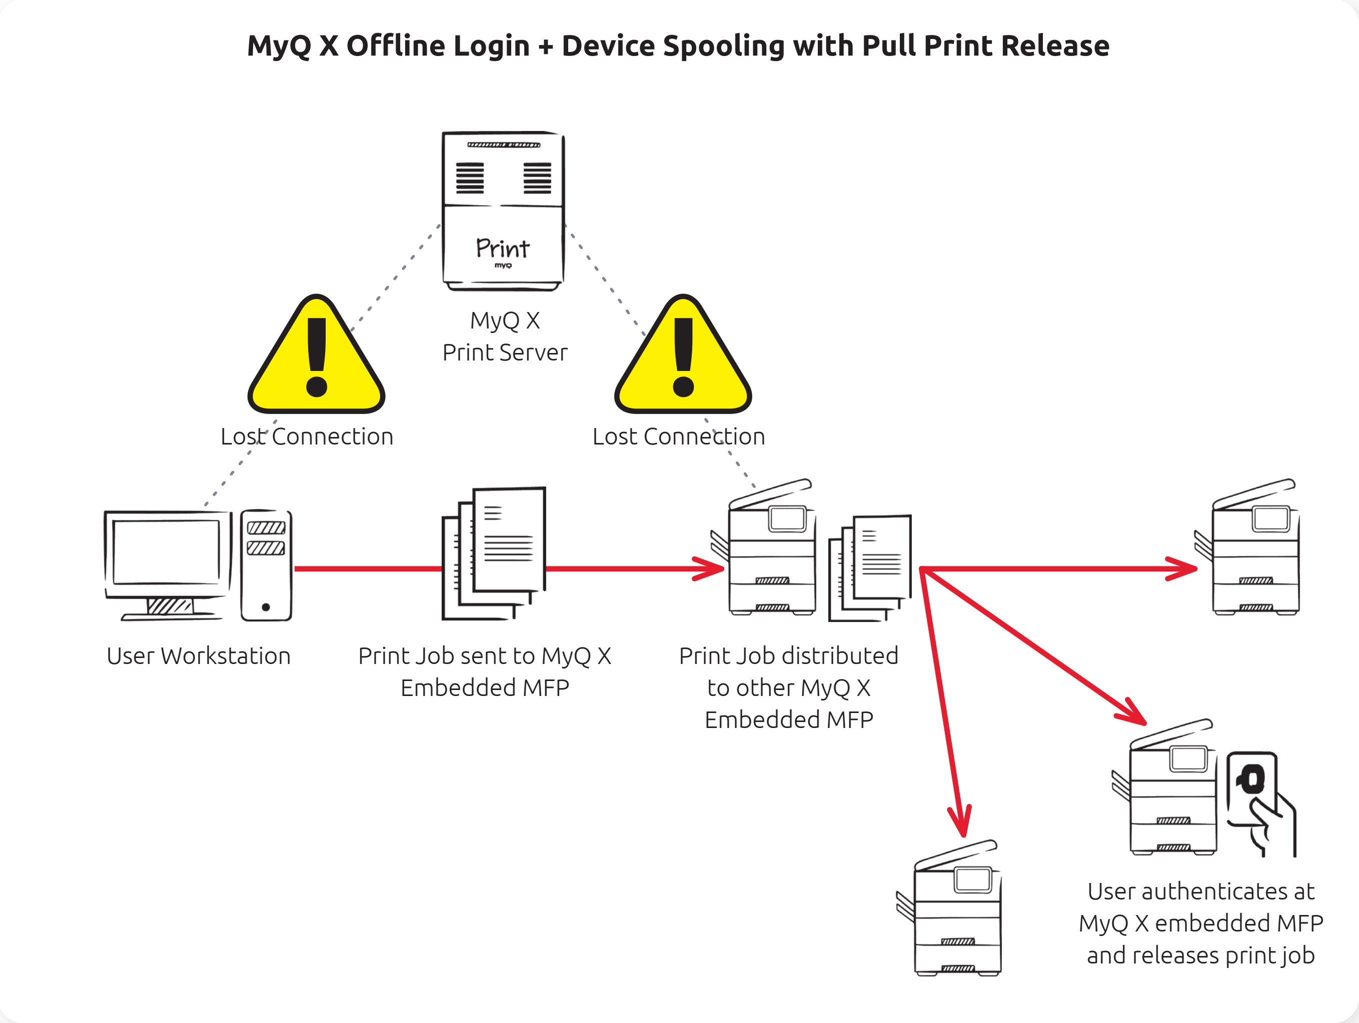

Figure 5.3: MyQ X Offline Login + Device Spooling with Pull Print

The **benefit of using MyQ X Device spooling and Offline login is that users will still be able to securely authenticate at the MFP, make copies and print with the convenience of print release** options such as Pull Print during plan or unplanned downtime of the server. No additional software client is required on the user's workstation. In addition, accounting of offline printing and copying will be recorded and updated to the MyQ X server when the connection is re-established.

Take note that services like scanning and advanced workflows which require an active server connection will not be available during downtime. MyQ X Offline Login and Device Spooling functionality is **only supported on selected**

vendors, as our MyQ X embedded terminal is frequently updated it is advised to check with your local MyQ partner for an updated list of supported devices.

5.4 Reduce Network and Server load with MyQ X Desktop Client and Client Spooling

With the MyQ X Desktop Client and the Client Spooling feature enabled, only job metadata is sent to the MyQ X server when the user prints a file. The metadata is extracted from the print job during the print action via MDC's local job parsing functionality on the users' workstation. **By processing print jobs locally on the user's workstation, it benefits the MyQ X Print server in terms of less compute resources are required** compared to parsing everyone's print jobs centrally at server level. After the print job has been parsed it will be stored on the user's workstation.

Job metadata only contains information about the print job (such as color, number of pages, job owner, etc.). It is, therefore, significantly smaller in data size compared to that of the entire print job. The MyQ X Print Server will receive the job metadata via MDC and register the print job information to the user's account.

Once the user authenticates at the MyQ X supported MFP (Multifunctional Printer), the MyQ X Print Server will share job data with the device. The user can then select to release the print job, which will prompt the MDC (MyQ Desktop Client) client to send the print job directly from the user's workstation to the relevant MFP. As a result, **MyQ X's Client Spooling feature can significantly decrease the server and network load** on the Print Server compared to traditional Printing methods.

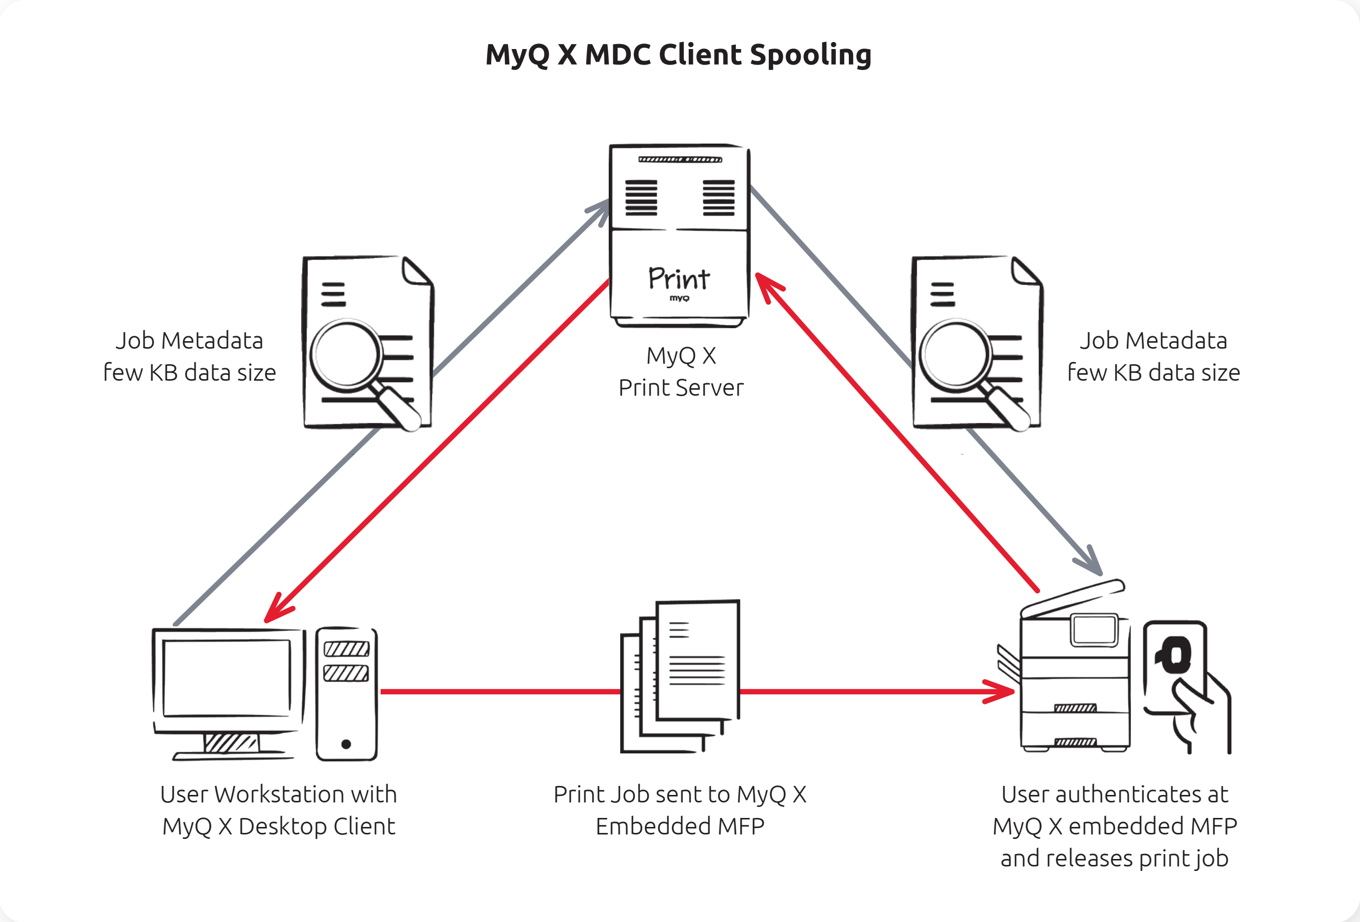

Figure 5.4: Reducing network and server load by using MDC Client Spooling

6. Disaster Recovery

6.1 What is IT Disaster Recovery?

Disaster recovery is a critical process designed to help organizations regain access and functionality to IT systems after a natural disaster, human error, or cyberattack. Similar in determining the right level of high availability, developing a Disaster Recovery Plan starts with conducting a BIA (Business Impact Analysis) to understand what impact a particular disaster will have on the business. The BIA determines the likelihood of potential risks, evaluates steps an organization can take to avoid or mitigate the risks, prioritizes responses, and estimates the monetary impact on the business.

6.2 How does disaster recovery work?

Disaster recovery relies upon the replication of data and computer processing in an off-premises location not affected by the disaster. When servers go down because of a natural disaster, equipment failure or cyber-attack, a business needs to recover lost data from a second location where the data is backed up. Ideally, an organization can transfer its computer processing to that remote location as well to continue operations.

6.3 The different types of IT disaster recovery

There are multiple methods by which an organization can protect and recover its data from a disaster. These methods can range from having data backups in off-site locations, virtualization, and Disaster Recovery as a Service (DRaaS) to having costly secondary hot sites with up-to-date data. Having secure **off-site backups might be sufficient for most organizations as it is the most cost-efficient.** However, there are some limitations in that it only covers data and does not provide replacement infrastructure if your primary site suffers from a disaster. Again, as an organization, you need to decide on a Disaster recovery plan that meets your business requirements and for which you can justify the cost.

6.4 Differences between High Availability and Disaster Recovery

High Availability

- High Availability minimizes downtime.
- Eliminating single points of failure is at the core of High Availability.
- High Availability helps mitigate the risk of hardware failure but **does not protect against data loss**.
- High Availability - Synchronous

Disaster Recovery

- Disaster recovery is at the center of dealing with worst-case scenarios and how to get your storage systems up as quickly as possible. DR protects you from situations that could otherwise be detrimental to your business.
- Geographically separated backups are at the center of disaster recovery.
- Disaster recovery is high-level in design and consists of a combination of a plan and technology design. High Availability is more about the technology design, combining failovers and redundancy to eliminate single points of failure.
- Disaster recovery – Asynchronous

7. Conclusion

As we discussed in the whitepaper MyQ X provides a set of free tools and features to aid organizations in keeping print highly available for end users in times of planned and unplanned downtime. However, we do understand that there are organizations to whom print management plays a critical role in their daily operations that require a high availability and disaster recovery strategy.

Determining the right level high availability for the infrastructure supporting your print environment is a calculation unique to each organization. **For most organizations the deciding factor comes down to the cost of implementing high availability and disaster recovery strategies versus the financial and business impact of having downtime in your IT system continuity.** Also considering that not all IT services are equally critical to business operations.

While high availability concepts are primarily technology-centric, disaster recovery encompasses more than just software and hardware elements. High availability focuses on addressing isolated failures in an IT system, while Disaster recovery deals with failures of a larger scope and the consequences of such failures.

High Availability on its own cannot ensure protection from disasters but can efficiently complement disaster recovery strategies. **To efficiently protect your data and ensure IT continuity, MyQ recommends that each organization should consider a high availability strategy that includes a Disaster Recovery plan for IT services. In addition to the high availability and disaster recovery, IT can use MyQ X's set of free tools to add an additional layer of failover protection.**