FINASTRA

FINANCE IS OPEN

Treasury and Capital Market

Finastra Kondor OpenShift Infrastructure Guide

Contents

Disclaimer

This document is written for Finastra Kondor 3.6.0 (Front Office, Back Office, and Risk Standard Limits & Market Risk) using Red Hat OpenShift 4.18. Its focus is on **Self-Managed On-Premises OpenShift deployment** and does not cover OpenShift deployments on Cloud infrastructure such as Microsoft Azure, Amazon AWS, Google Compute Platform, or Oracle OCI.

Introduction

Confronted with the need to broaden their mandate beyond the traditional role of funding and monitoring investment risk, Organization treasuries are increasingly being called up to manage capital and to adapt to new regulatory and market pressures.

In this context, upgrading outmoded infrastructure and leveraging new technologies for their treasury systems becomes a priority. Such digital transformation requires scaling IT to support rapid delivery of new applications and services in a hybrid cloud environment, where applications and processes can take advantage of current technologies such as containers and microservices.

The portability and repeatability of containers can create cost and resource savings, coupled with a faster time to market and rapid innovation. Containers come with small overhead costs, helping to lower hardware, maintenance, and licensing costs. They can be implemented quickly, while components can be shared among containers.

To help financial institutions modernize their treasury platform and start taking advantage of containers at optimal costs, Finastra and Red Hat have collaborated on an end-to-end platform architecture with different infrastructure requirements for running Kondor Treasury System on Red Hat OpenShift Container Platform.

Available in several cluster sizes, this solution is designed for the hybrid cloud and is also customizable and fully interoperable with existing infrastructure. Enterprises can use this solution to start driving digital transformation through agile DevOps methodologies and quickly release new services with efficiency and scalability. The reference architecture detailed in this document can guide you through the often-complex hardware and software selection process when setting up the container platform and planning for the deployment of relevant software components for Fusion Kondor.

This document provides details for three baseline configurations based on business volume drivers:

F2B

Configuration	Users	Live Transactions	Daily Transactions
Small	10	5000	<=2000
Medium	50	50 000	5000
Large	>=100	1 000 000	>=5000

Risk

Configuration	Users	Live Transactions	Daily Transactions	# Limits
Small	10	10 000	<=100	1000
Medium	25	50 000	<=1000	5000
Large	30	100 000	>1000	20 000

Infrastructure Requirements Overview

Regardless of whether you are new to the Red Hat OpenShift platform or if you are an existing user, deploying Finastra Kondor requires certain **mandatory** infrastructure components, as well as certain **optional** requirements.

For each of these mandatory and optional requirements, several alternatives are available that the Organization will need to decide on. These requirements can be provided by either one of the different editions of the OpenShift platform (OKE, OCP, OPP) or by third-party components. In many cases, the Organization will already have some of these components deployed inside their existing infrastructure and can leverage these components for Finastra Kondor on Red Hat OpenShift.

Mandatory Infrastructure Requirements

Mandatory requirements consist of components and services that are **technically required** for running Finastra Kondor on Red Hat OpenShift. Many of these requirements are fulfilled by one of the three available editions of Red Hat OpenShift (OKE, OCP, OPP), but they can also be fulfilled by existing third-party solutions within their infrastructure.

The section below presents the mandatory infrastructure requirements:

- **An OpenShift Kubernetes** Cluster, which consists of:
  - **Master Nodes** or the **Control Plane** - responsible for cluster management and for providing the API that is used to configure and manage resources within the Kubernetes cluster.
  - **Worker Nodes** or the **Data Plane** - responsible for running the user workloads, i.e., Finastra Kondor.
  - **Infrastructure Nodes** - regular worker nodes that are specifically designated to run OpenShift's router, image registry, and other infrastructure services.
- **Router / Ingress** - responsible for handling incoming HTTP/HTTPS requests and dispatching them for processing to the specific services that handle them.
- **Container Image Registry** - a repository for Container Images downloaded onto each Worker Node. OpenShift requires one of the following:
  - **Basic Container Image Registry**
  - **Enterprise-grade Registry** - provides enhanced functionality over a Basic Kubernetes Container Image Registry. Typically, an enterprise repository manager offers features such as automatic mirroring of external repositories, the ability to act as a repository for other types of binaries like Maven artifacts, automatic security scanning, and more.
- **Load Balancer** - a device or application that acts as a single network access point to a Kubernetes cluster. Applications outside of the Kubernetes Cluster connect to the Load Balancer, which distributes the traffic across the different Worker (or Master) Nodes. In the context of Finastra Kondor specifically, the Load Balancer is required to support both ISO/OSI standard Layer 7 (for HTTP/HTTPS traffic) *and* Layer 4 (TCP traffic) Load Balancing.
- **Storage** - Finastra Kondor requires persistent storage to be added to the OpenShift cluster. Typically, this involves an *external* storage solution that has a supported Kubernetes CSI driver that meets the following requirements:
  - POSIX compliant
  - Persistent
  - Read/Write Multiple Pods
  - Dynamic Provisioning.
- **DNS (Domain Name System)** - Finastra Kondor requires access to the Organization's DNS infrastructure. It is assumed that DNS is available *by default* in any Organization.
- **Database** - Finastra Kondor requires access to a relational database, specifically either Microsoft SQL Server (2019, 2022) or SAP Sybase ASE (16 SP03, 16 SP04; check our compatibility matrix for specific certified Patch Levels). This database typically runs *outside* the OpenShift cluster on dedicated Virtual Machines or Bare Metal servers.
- **Standalone (Virtual) Machine(s)** - although Finastra Kondor can be fully deployed inside an OpenShift cluster, it depends on several components typically deployed on one or more Standalone Virtual Machines or Bare Metal servers. Examples of such components include:
  - **Finastra Entitlement Manager (EM, former ELS)** - *must* be deployed on Standalone (Virtual) Machines, one for the primary and one for the secondary EM instance
  - **Custom Interfaces** – along with the batch jobs, they are typically deployed on Standalone Virtual Machines.

Note: These components can usually be deployed **on the same** (Virtual) Machine(s).

Optional Infrastructure Requirements

Optional requirements consist of components that are not technically required to deploy and run Finastra Kondor on Red Hat OpenShift. However, many Organizations may consider them desirable or even mandatory from a policy/compliance point of view. Often, Organizations already have existing third-party solutions within their infrastructure that are compatible with OpenShift.

The section below presents the optional infrastructure requirements:

- **Multi-Cluster Management** – for Organizations operating multiple OpenShift clusters, Multi-Cluster Management is a compelling addition to streamline the management of all their OpenShift clusters from one single point.
- **Metrics/Usage Monitoring -** a solution that can monitor anything happening inside the cluster that can be expressed as a number such as:
    - Resource availability (in case the required number of instances of a specific deployment/pod is currently running)
    - Pod Health Status
    - CPU Utilization
    - Memory Utilization
    - Disk I/O
    - Number of API calls.

  Typically, the Metrics/Usage Monitoring can also raise alerts in case specific "metrics" exceed or fail to meet a specified threshold.
- **Log Monitoring** - a solution that can monitor the log files generated by Pods inside the cluster and scan for specific patterns in those log files indicating abnormal behavior. This component can also raise alerts when such abnormal patterns are detected, maintain a history of log files, and provide easy and fast ways for administrators to manually search and analyze log files to support incident investigation.
- **Standalone (Virtual) Machine(s)**
    - **Administrator Deal Manager instance** - it is convenient to have a Red Hat Enterprise Linux (Virtual) Machine server available where technical administrators can run Kondor Deal Manager.

    Note: If Kondor Vision is included, the Administrator Deal Manager instance is completely optional, since it can be provided through "Kondor Classic Viewer" hosted inside the OpenShift cluster.

    Note: This (Virtual) Machine requirement can usually be **combined** with the (Virtual) Machine(s) for the Finastra Entitlement Manager and Custom Interfaces.

Detailed Requirements and Available Implementations

This section provides further details on the mandatory and optional requirements and discusses the different options available to Organizations to fulfill those requirements.

Mandatory Infrastructure Requirements

OpenShift Kubernetes Cluster

There are three Red Hat OpenShift editions available. Each edition corresponds to a specific *Red Hat support scope*. For a comparison of features between each edition, please see Appendix 1 of the Self-Managed Red Hat OpenShift subscription guide.

OpenShift Kubernetes Engine (OKE) is the base-level OpenShift edition, which includes Red Hat support for all the mandatory OpenShift features.

OpenShift Compute Platform (OCP) adds on top of OpenShift Kubernetes Engine to provide more developer-related features. In the context of Finastra Kondor, it adds Red Hat support for:

- User Application and OpenShift platform logging based on the OpenShift "Loki Operator"
- User Workload Monitoring (metrics)
- Service Mesh, which can be optionally enabled for Finastra Kondor.

OpenShift Platform Plus (OPP) adds on top of OpenShift Compute Platform to provide Red Hat support for enhanced management and security features, as well as the Quay Global Registry product. For Organizations managing multiple OpenShift Clusters, the Multi-Cluster Management features can significantly streamline their cluster management operations.

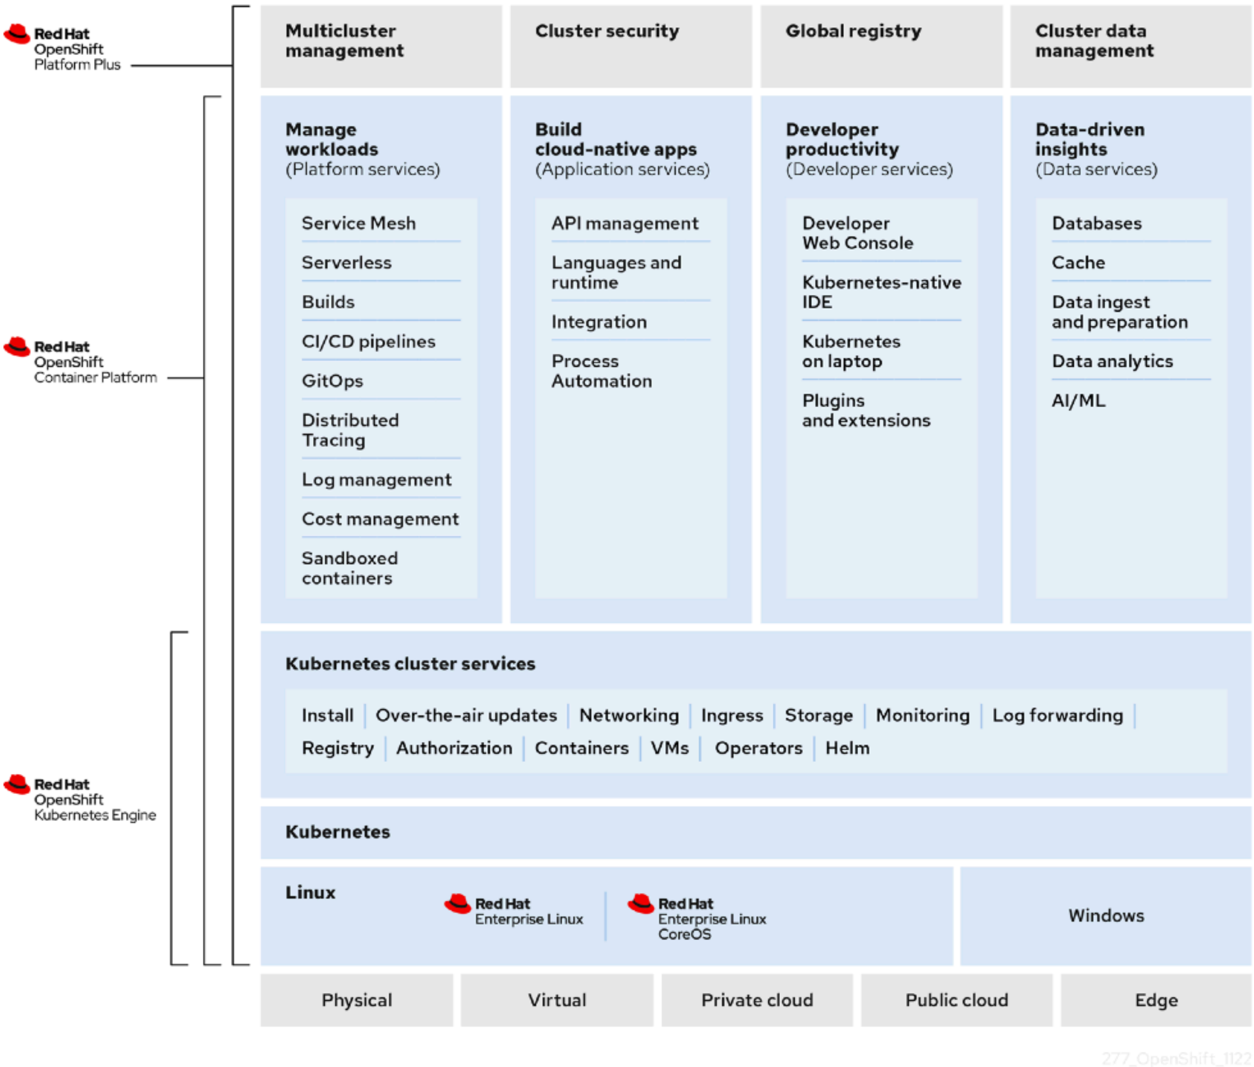
Red Hat OpenShift Platform Plus
Red Hat OpenShift Container Platform
Red Hat OpenShift Kubernetes Engine
Multicluster management
Cluster security
Global registry
Cluster data management
Manage workloads (Platform services)
Service Mesh
Serverless
Builds
CI/CD pipelines
GitOps
Distributed Tracing
Log management
Cost management
Sandboxed containers
Build cloud-native apps (Application services)
API management
Languages and runtime
Integration
Process Automation
Developer productivity (Developer services)
Developer Web Console
Kubernetes-native IDE
Kubernetes on laptop
Plugins and extensions
Data-driven insights (Data services)
Databases
Cache
Data ingest and preparation
Data analytics
AI/ML
Kubernetes cluster services
Install | Over-the-air updates | Networking | Ingress | Storage | Monitoring | Log forwarding
Registry | Authorization | Containers | VMs | Operators | Helm
Kubernetes
Linux
Red Hat Enterprise Linux
Red Hat Enterprise Linux CoreOS
Windows
Physical
Virtual
Private cloud
Public cloud
Edge
277_OpenShift_1122

Choosing the right OpenShift edition for your Organization depends on several factors such as objectives and existing infrastructure. This document aims to provide support in making the optimal decision.

You can extend the OpenShift capabilities by using Operators. Operators can provide extended functionality or connect OpenShift with third-party solutions such as Load Balancers, Log/Metrics monitoring solutions, and more.

You can find the list of available operators organized by category on the OpenShift Operator Hub, accessible via each OpenShift Cluster, and on https://catalog.redhat.com/.

Red Hat-supported operators display a "Red Hat badge", and the support agreement linked to your OpenShift edition will determine if you are entitled to a certain operator. Community Operators are not supported by Red Hat, and Operators provided by Red Hat partners are supported by the partner and have a different lifecycle. Please follow the guidelines when integrating these.

OpenShift Router/Ingress

Every Kubernetes Cluster requires a way to enable external traffic to enter the cluster in a controlled manner. In the Kubernetes context, this is handled by an Ingress Controller, which OpenShift *extends upon* with its Router concept.

There are currently two possible options for a Router in OpenShift:

- OpenShift default HAProxy Router
- F5 BIG-IP Container Ingress Services.

Note: The use of F5 BIG-IP Container Ingress Services requires a separate license from F5.

Many organizations have standardized on F5 BIG-IP as their Enterprise-grade Load Balancer; therefore, they will decide to use F5 BIG-IP Container Ingress Services. At the same time, this would also fulfill the requirement for a Load Balancer (for more information, see the "Load Balancer" section below).

Organizations that choose to use the OpenShift default HAProxy router need to also select a suitable Load Balancer.

Basic Container Image Registry

A Basic Container Image Registry is included in every edition of OpenShift. While it is not recommended for high-volume traffic, it is completely suitable for the low-volume requirements of Finastra Kondor on the registry.

When deploying Finastra Kondor into an OpenShift cluster, the Worker Nodes need to *pull* (download) the relevant *container images* from a *container registry*. Finastra makes available its Helm charts and Container Images related to Finastra Kondor through its public *registry*. Most financial Organizations do not allow access to external registries from inside their Data Center. Therefore, the Finastra container images need to be *mirrored* to a local container registry.

The Basic container registry included in OpenShift does not support automatic mirroring of external registries. Finastra makes available a set of Shell Scripts that *manually* perform the mirroring task by first pulling all required images from their public container registry, and then *pushing* them into the local Basic Container Registry.

The Basic Container Registry also does not provide any functionality beyond *pull* and *push*. If an Organization needs more advanced functionality, it should consider including an Enterprise-Grade Registry in the solution.

Enterprise-Grade Registry

Enterprise-Grade Registry solutions offer many features over the Basic OpenShift Container Registry:

- automatic repository/registry mirroring
- high availability
- superior performance
- enhanced security, authentication, and permissions
- support for other types of package formats such as Maven artifacts, NPM packages, NuGet packages, etc.
- automatic artifact scanning/analysis.

While the Enterprise-Grade Registry comes with a higher complexity, it may be the right choice for an Organization seeking a binary repository manager for their IT department.

Examples of binary repository managers include:

- Red Hat Quay (included in OpenShift Platform Plus, as well as available separately)
- JFrog Artifactory
- Sonatype Nexus.

If your Organization already has one of these binary repository managers inside its existing infrastructure, it can be directly used by Finastra Kondor on OpenShift.

Load Balancer

Load Balancer serves the following purposes:

- provides a **single Highly Available IP address** for access to the cluster
- distributes Kubernetes API calls across each Master Node
- distributes HTTP(S) requests across the Router instances on each Infrastructure Node
- forwards **raw TCP/IP traffic on specific ports** to services on the different Worker Nodes.

The distribution of Kubernetes API calls and HTTP(S) requests relies on ISO/OSI Model Layer 7 Load Balancing, while the handling of raw TCP/IP traffic relies on ISO/OSI Model Layer 4 Load Balancing.

Although the Load Balancer logically exists *outside* the OpenShift cluster, it is nevertheless tightly coupled with OpenShift. Therefore, not all Load Balancers are suitable.

There are currently two suitable Load Balancer options for *a Self-Managed On-Premises OpenShift Cluster*:

1. MetalLB, included with OpenShift as an Operator
2. F5 BIG-IP Container Ingress Services.

Note: Using F5 BIG-IP Container Ingress Services requires a separate license from F5.

Storage

The applications running inside the OpenShift cluster need access to persistent storage. OpenShift (or Kubernetes) does not provide a “one-size-fits-all” storage solution. Instead, Kubernetes includes a *standard*: “Container Storage Interface” (CSI) standard. Of this standard, there can be many different implementations, called “CSI drivers”.

There are currently approximately 141 different CSI drivers available for Kubernetes (and therefore OpenShift) related to different underlying storage solutions. In theory, an Organization can select any one of these CSI drivers, as long as they meet the following list of requirements:

- POSIX compliant
- Persistent
- Read/Write Multiple Pods
- Dynamic Provisioning
- The **underlying physical storage** supports data replication between the Primary Data Center and the Disaster Recovery Data Center.

Finastra confirms compatibility with the following CSI drivers/underlying storage solutions:

- NFS – It is strongly recommended that the NFS *server* be Highly Available.
- CephFS – The underlying file system can be provided by OpenShift Data Foundation or by another Ceph-based storage solution.
- Proxworx – Requires the use of the Sharedv4 PVCs feature.

Other CSI drivers may be suitable as well. If you want to use a CSI driver that is not mentioned in the list above, please reach out to your Finastra Customer Advocate.

Domain Name System (DNS)

Finastra Kondor creates many hostname-based Routes:

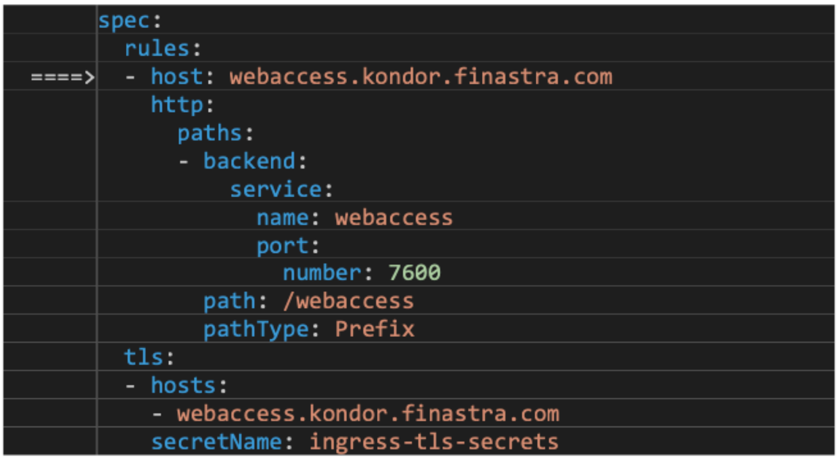

```
        spec:
          rules:
====>     - host: webaccess.kondor.finastra.com
            http:
              paths:
              - backend:
                  service:
                    name: webaccess
                    port:
                      number: 7600
                path: /webaccess
                pathType: Prefix
          tls:
          - hosts:
            - webaccess.kondor.finastra.com
            secretName: ingress-tls-secrets
```

Each of these host names *must* be resolvable via DNS for external clients. To avoid having to create multiple DNS “Address records”, it is strongly recommended to create a single CNAME record pointing to the “Address record” for Load Balancer:

```
*.kondor.finastra.com.  3600 IN CNAME  lb.kondor.finastra.com.
lb.kondor.finastra.com. 3600 IN A      10.204.16.129
```

Optional Infrastructure Requirements

Multi-Cluster Management

Red Hat Advanced Cluster Management for Kubernetes, which is included in Red Hat OpenShift Platform Plus, controls multiple clusters and applications from a single, central console, with built-in security policies to offer compliance and maintain consistency. This is helpful if you want to add more features to the Red Hat OpenShift by deploying apps, managing multiple clusters, and enforcing policies across multiple clusters at scale.

When it comes to advanced security, Red Hat Advanced Cluster Security for Kubernetes (ACS) is a security-focused solution designed to protect containerized workloads running on Kubernetes platforms, including Red Hat OpenShift. While Advanced Cluster Management (ACM) is primarily for multi-cluster management and policy enforcement, ACS specializes in security for the entire application lifecycle within and across OpenShift clusters and is included in Red Hat OpenShift Platform Plus.

Metric/Usage Monitoring

Although it is an optional component, most Organizations will want or need to monitor and report on the availability and utilization of both Finastra Kondor (the "User Workload") and the OpenShift Cluster (master nodes, worker nodes, core Kubernetes components, infrastructure components).

All OpenShift editions include a monitoring solution based on Prometheus and Grafana. However, in OpenShift Kubernetes Engine, this solution is only allowed to be used to monitor the OpenShift Cluster itself, *not User Workloads*. When it comes to OpenShift Compute Platform and OpenShift Platform Plus, this monitoring solution can be used for free to monitor User Workloads as well.

Organizations can use other monitoring solutions with OpenShift, and many organizations do. Below are some examples of why an Organization may use a third-party monitoring solution:

- You have an existing monitoring solution that is used as standard across your organization.
- You have advanced monitoring requirements that cannot be provided by OpenShift Prometheus/Grafana.

Examples of monitoring solutions compatible with OpenShift include:

- DynaTrace
- Prometheus/Grafana.

For other available metrics monitoring solutions, please check https://catalog.redhat/com or the OpenShift Operator Hub.

In case you already have an existing monitoring solution, please reach out to the respective vendor for information on compatibility with OpenShift.

Log Monitoring

Most Organizations will want or need to collect and analyze log files from both Finastra Kondor (the "User Workload") and the OpenShift Cluster.

All OpenShift editions allow the installation of the "Loki operator". This is a horizontally scalable, highly available, multi-tenant log aggregation system. You can use it to collect and aggregate logs from the OpenShift cluster itself as well as from User Workloads. Logs can be searched, and Alerting Rules can be defined on specific log messages as well as on items like numbers of specific messages over a period of time, such as "Max API call rate exceeded". Loki log *storage* is optimized for *short-term storage*, typically in the order of 7 days.

Organizations are free to use other log aggregators with OpenShift. Below are some examples of why an Organization may use a third-party log collection & analysis solution:

- You have an existing log monitoring solution that is used as standard across your Organization.
- You have more advanced requirements that cannot be met by Loki.
- You are required to keep a log file history for longer than the Loki log storage was designed for.

Examples of log aggregators compatible and popular for use with OpenShift include:

- Splunk
- DynaTrace
- ELK / Elastic Stack.

For other available log monitoring solutions, please check https://catalog.redhat/com or the OpenShift Operator Hub.

In case you already have an existing monitoring solution, please reach out to the respective vendor for information on compatibility with OpenShift.

Summary of Requirements and Available Implementations

Infrastructure Requirement	OKE	OCP	OPP	Third-Party
Router/Ingress	Yes, HAProxy-based			F5 BIG-IP
Basic Container Image Registry	Yes, Kubernetes Image Registry			N/A
Load Balancer	Yes, MetalLB			F5 BIG-IP
Storage	No		Yes, ODF/Ceph FS	NFS CephFS Portworx <others>
DNS	N/A, assumed available by default			
Multi-Cluster Management	No		Yes	N/A
Enterprise-Grade Registry	No		Yes, Quay	JFrog Artifactory Sonatype Nexus <others>
Metrics/Usage Monitoring	Yes, Prometheus/Grafana			DynaTrace Prometheus/Grafana <others>
Log Monitoring	Yes, Loki			Splunk ELK / Elastic Stack <others>

Guided OpenShift Infrastructure Selection

This section is aimed at Organizations that use Red Hat OpenShift for the first time. Those that do already have OpenShift as part of their IT infrastructure are likely to already have solutions for the various functions listed below and should leverage those existing implemented solutions when deploying Finastra Kondor on OpenShift.

Database Selection

Finastra recommends using a supported version of Microsoft SQL Server on Windows Server. If you are an existing Finastra Kondor customer using SAP Sybase ASE, we recommend migrating to Microsoft SQL Server. If not possible, stick with your existing SAP Sybase ASE.

OpenShift Edition Selection

The following questions can help you choose the right OpenShift edition for Finastra Kondor:

- Do I need Multi-Cluster Management?
- Do I need one of the OpenShift Logging or Monitoring solutions?

If you are a current user of Finastra Kondor doing a “like for like migration” and you *do not currently* monitor Kondor usage and logs, then you *do not* need to add OpenShift monitoring solutions, and you *do not* need OpenShift Compute Platform.

If you already have existing metrics/logs monitoring solutions within your IT infrastructure, there is a good chance you can use this type of solution with Finastra Kondor on OpenShift, and you *do not* need OpenShift Compute Platform.

Metric & log monitoring solutions require *extensive* additional compute and storage resources. If you do add such solutions to the scope of a Finastra Kondor OpenShift deployment, make sure to separate the cost of such solutions from the cost of “core Finastra Kondor”. See the Deployment Sizing section from the Logging using LokiStack for details on the resource requirements of OpenShift LokiStack.

For more information, please see the *Metrics/Usage Monitoring Selection* and *Log Monitoring Selection* sections below.

Note: Use **OpenShift Platform Plus** if you need Multi-Cluster Management. Use **OpenShift Compute Platform** if you need the **OpenShift Prometheus/Grafana and/or LokiStack** solutions to monitor the Finastra Kondor workload. Otherwise, use **OpenShift Kubernetes Engine**.

It is recommended for existing users of Red Hat OpenShift to remain on their current OpenShift edition.

Router & Load Balancer Selection

If you use F5 BIG-IP, we recommend you consider acquiring the required licenses for the F5 BIG-I Container Ingress Service, and:

- Use F5 BIG-IP Container Ingress Service for Ingress/Router.
- Use F5 BIG-IP Container Ingress Service as Layer 4 Load Balancer.
- Use F5 BIG-IP as Layer 7 Load Balancer.

If you have F5 BIG-IP but cannot acquire the licenses for the F5 BIG-IP Container Ingress Service:

- Use the default OpenShift HAProxy for Ingress/Router.
- Use MetalLB as Layer 4 Load Balancer.
- Use F5 BIG-IP as Layer 7 Load Balancer.

If you have another Enterprise-Grade Load Balancer:

- Use the default OpenShift HAProxy for Ingress/Router.
- Use MetalLB as Layer 4 Load Balancer.
- Use your existing Enterprise-Grade Load Balancer as Layer 7 Load Balancer.

If you do *not* have an existing Enterprise-grade Load Balancer:

- Use the default OpenShift HAProxy for Ingress/Router.
- Use MetalLB as Layer 4 Load Balancer.
- Use MetalLB as Layer 7 Load Balancer.

Storage Selection

Most Organizations have access to an Enterprise-Grade Highly Available storage solution, typically in the form of a SAN or a NAS. If there is a CSI Driver that is compatible with Finastra Kondor requirements, as well as this existing storage solution, this existing storage solution can likely be used for Finastra Kondor on OpenShift.

In case your Organization does *not* have an existing Enterprise-grade Highly Available storage solution, please reach out to your Red Hat Customer Advocate to discuss the potential solutions.

Enterprise-Grade Registry Selection

If you have an existing Enterprise-Grade Container Registry like JFrog Artifactory or Sonatype Nexus, use your existing registry.

If you do *not* have an existing Enterprise-Grade Container Registry and you need *automatic* mirroring of the Finastra Registry or other advanced features, then:

- Use Red Hat Quay if you have selected OpenShift Platform Plus.
- If you have not selected OpenShift Platform Plus, acquire one of the following:
  - Red Hat Quay
  - JFrog Artifactory
  - Sonatype Nexus.

If you do not need automatic mirroring of the public Finastra Container Registry or any other advanced features, use the default internal Container Registry.

Metrics/Usage Monitoring Selection

In case you have an existing metrics/usage monitoring solution, reach out to your existing Vendor to check if it supports monitoring of the OpenShift Platform. If it does, we strongly recommend using your existing solution, even if this requires acquiring additional licenses.

If you do *not* have an existing monitoring solution, and you selected one of the OpenShift Compute Platform or OpenShift Platform Plus editions, use the standard Prometheus/Grafana monitoring included in these editions.

In case you selected the OpenShift Kubernetes Engine edition, you need to decide whether you only need to monitor the OpenShift infrastructure (Master Nodes, Infrastructure Nodes, Worker Nodes, but *not* the user workload), or if you also need to monitor the user workload (e.g. individual Finastra Kondor components).

- If you need to only monitor the OpenShift infrastructure, then use the standard Prometheus/Grafana monitoring that is included in OpenShift Kubernetes Engine.
- If you need to monitor the Finastra Kondor User Workload as well, choose one of the various third-party monitoring solutions that are available.

Log Monitoring Selection

If you have an existing log collection and monitoring solution, reach out to your existing Vendor to check if it can consume logs from the OpenShift Platform and User Workloads. If it does, we strongly recommend using your existing solution.

If you do *not* have any existing log collection and monitoring solution, you need to decide if your new Finastra Kondor deployment requires log file monitoring. If the answer is “Yes”, we suggest the following options:

1. Use the “Loki” solution that comes with Red Hat OpenShift. The downside of this is that the *Log Storage* of Loki is optimized for short-term storage only, and only a short log file history will be kept.
2. Build your own log collection and monitoring solution using solutions like the ELK Stack.
3. Acquire a third-party solution.

If you are migrating from a “legacy” Finastra Kondor deployment to Finastra Kondor on OpenShift *and you did not monitor log files* in your legacy environment, you probably do not need a log collection and monitoring solution with Finastra Kondor on OpenShift.

Finastra Kondor T-Shirt Sizing

This section presents the indicative compute and the memory requirements for the **Data Plane** based on the T-Shirt definitions provided in the *Introduction* section. In addition to the Data Plane, a full solution needs to include at least:

1. a database
2. a standalone VM hosting the Finastra License Server (EM) and possibly interfaces, and other non-containerized applications.

For High Availability purposes, the number of databases and standalone VMs may need to be doubled.

Please keep in mind that these are just *indications*. Each client is different, and a T-Shirt sizing can never capture all characteristics that influence capacity planning. We strongly recommend reaching out to your Finastra Customer Advocate to request a *bespoke* sizing for *your* specific situation. In addition, especially if you are new to the Red Hat OpenShift ecosystem, we recommend reaching out to your Red Hat Customer Advocate for specific advice on Control Plane capacity and Infrastructure requirements.

The below estimations *do not include redundancy* (for more information, please see the *Adding redundancy to T-Shirt estimates* section below). If Vision components are part of the solution, please reach out to your Finastra Kondor Customer Advocate.

Worker Nodes need to have *100 GB of storage* to hold the locally cached container images and to accommodate deployment and upgrades of Finastra Kondor.

F2B (without Vision)

Configuration	vCPU	RAM
Small	16	64GB

Medium	24	96GB
Large	36	128GB

Kondor Classic Viewer (Kondor QT client running inside worker nodes)

In case the Kondor QT clients are not operated inside the worker nodes, an equivalent sizing for VMs needs to be provisioned.

Configuration	vCPU	RAM
Small	10	24GB
Medium	50	128GB
Large	100	256GB

Note: Resources should be scaled to accommodate the cases where the number of users is >100.

Risk Standard Limits & MRP

For Standard Limits:

Configuration	vCPU	RAM
Small	6	19GB
Medium	6	19GB
Large	9	28GB

For Market Risk:
Resource sizing for Market Risk involves many different input parameters including:

- the number of live/outstanding deals, *per type of deal*
- the types of Market Risk processes, such as Historical VaR, Monte-Carlo VaR, PFE/CVA, Sensitivity VaR, Stress Testing, Back Testing, as well as the *number* of each of those processes
- the configuration of each of these processes, such as the number of Risk Classes/Factors to compute, the number of scenarios to apply, the days horizon, etc.

The exact process of sizing to account for Market Risk goes beyond the scope of this document. Please keep in mind that if you require Market Risk as part of your solution, additional compute and memory resources need to be added over and above those required for “Standard Limits”.

Adding Redundancy to T-Shirt Estimates

A typical OpenShift Cluster consists of:

- three Master Nodes (Control Plane)
- three Infrastructure Nodes
- *at least* three Worker Nodes (Data Plane).

Each of these nodes is expected to be *physically* deployed in *separate physical hardware*. This protects the Cluster from a physical server failure.

If there are three worker nodes on three physical servers and one server fails, the Cluster compute capacity is reduced to 2/3 of the total compute capacity. Our estimates are provided on the assumption that “required capacity == reduced capacity”.

Example 1

Number of Worker Nodes = 3
Required Capacity = 16 vCPU and 64 GB RAM
Required Capacity = Reduced Capacity = 2/3 Total Capacity
Total Capacity = Required Capacity * 3/2 = 16 / (2/3) = 24 vCPU and 96 GB RAM
Capacity per Worker Node = Total Capacity / 3 = 8 vCPU and 32 GB RAM

The $2/3$ factor comes directly from the number of physical servers over which the worker nodes are deployed.

Example 2

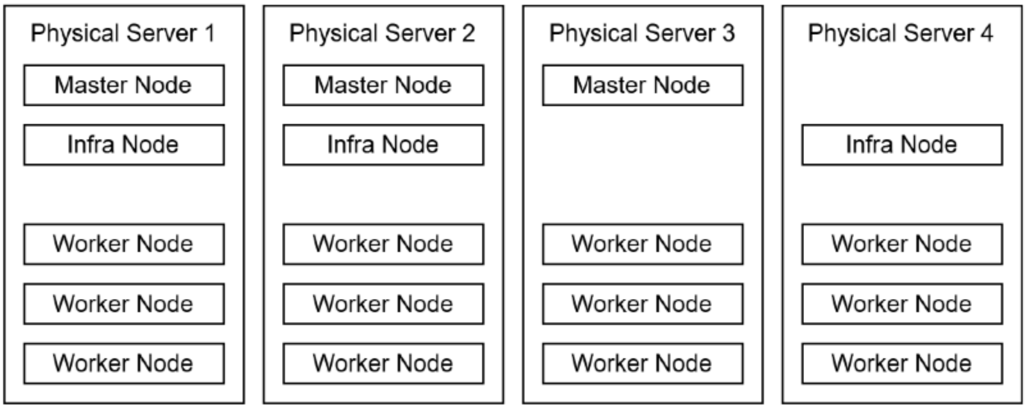

Figure 1: Four-server example

In this case, the client has four physical servers, hosting three Master Nodes, three Infrastructure Nodes, and **twelve** Worker Nodes. If one physical server fails, the client loses $1/4$ of the compute capacity, leaving $3/4$ of the total capacity of the Data Plane. This means that, to go from "required capacity" to "total capacity", the client needs to divide the required capacity by $3/4$.

Assuming an even distribution of X Worker Nodes over N Physical Servers, the formula used to get from "required capacity" $C_{required}$ to "total capacity" C_{total} is the following:

$$C_{total} = \frac{C_{required}}{1 - \frac{1}{N}}$$

Example 3

If the *required capacity* is 48 vCPU and 128GB RAM, and you want to distribute six Worker Nodes over six Physical Servers, then the Total Data Plane Capacity should be:

$vCPU_{total} = \frac{48}{1 - \frac{1}{6}} = 57.8$	$RAM_{total} = \frac{128}{1 - \frac{1}{6}} = 153.6$

Each worker node should therefore be sized with 10 vCPUs and 26 GB RAM (rounded up).

Building Examples

Example 1

Client overview

An existing Finastra Kondor user wants to migrate to an OpenShift-based deployment. The Organization is using Front Office, Back Office, Limits (Credit & Settlement Limits), and Market Risk (Historical VaR, Stress Testing, Back Testing). They will not use Kondor Vision.

The client's usage patterns are the following:

Measure	Value
Concurrent Front Office Users	16
Concurrent Back Office Users	24
Concurrent Limit Users	36
Average Trades/Day	130
Live Trades	765
Number of Credit/Settlement Limits	<1000
Market Risk	
HS VaR	8 processes, 261 days, 3 risk classes, 250 days results archived
Stress Tests	10 stress test scenarios, 250 days results archived
Back Tests	4 Back Test scenarios, 250 days results archived

In addition, the hypothetical Organization has the following existing infrastructure available:

- Microsoft SQL Server Database Servers using AlwaysOn
- Citrix infrastructure for serving DealManager from a central location
- Enterprise-Grade storage supporting NFS
- F5 BIG-IP Load Balancer.

The Organization agreed to acquire F5 BIG-IP Container Ingress Services so they can use their existing F5 BIG-IP Load Balancer for Ingress, both for Layer 7 HTTP(S) and for Layer 4 TCP traffic.

The Organization is not currently doing any metrics monitoring on the User Workload and is not monitoring any log files.

The Organization does not have an Enterprise-Grade Container Registry solution and is not planning to invest in such a solution. They are happy with using the included Basic Container Registry and shell scripts to periodically mirror the contents from the Finastra registry.

Because this client is new to OpenShift, they will need to deploy a Control Plane, a Data Plane, and Infrastructure Nodes. Because they will only deploy three clusters (one for Production, one for UAT and Dev, and one for DR), they do not require advanced cluster management. Because they do not need to monitor the User Workload, they do not need the OpenShift Compute Platform.

Therefore, the client should select **OpenShift Kubernetes Engine**.

T-Shirt size selection

The Organization falls into the "Small" category of T-Shirt size for both Finastra Kondor F2B and Risk. Therefore, they need the following Kubernetes compute capacities:

	vCPU	RAM
Kondor F2B	16	64
Risk – Standard Limits	6	19
Risk – Market Risk	4	8
Total	**26**	**91**

This sizing applies to the Production and DR environments. Since the T-Shirt sizing already falls into the smallest bracket, this same sizing applies to the UAT and Dev environments.

Note: As previously mentioned, the sizing for Market Risk requires additional compute and memory resources.

Sizing of Control Plane

The T-Shirt sizing of the Organization falls into the minimum recommended Control Plane sizing from Red Hat (see Control plane node sizing). Therefore, the Organization should deploy three Master Nodes with the following configuration:

	vCPU	RAM	Storage
Master Node 1	4	16	128 GB
Master Node 2	4	16	128 GB
Master Node 3	4	16	128 GB

Note: This configuration can be used for all three planned OpenShift clusters (Production, UAT+Dev, and DR).

Sizing of Infrastructure Nodes

The only infrastructure services this hypothetical Organization needs to run on the Infrastructure Nodes are the F5 BIG-IP Container Ingress Services, the Basic Container Registry, and Prometheus/Grafana (for cluster monitoring only). Therefore, the Organization can deploy minimally sized Infrastructure Nodes. For more details, please see Infrastructure node sizing.

Note: The Container Registry requires additional storage to physically store the container images.

	vCPU	RAM	Storage
Infrastructure Node 1	4	24	512GB
Infrastructure Node 2	4	24	512GB
Infrastructure Node 3	4	24	512GB

Note: This configuration can be used for all three planned OpenShift clusters (Production, UAT+Dev, and DR).

Sizing of the Data Plane

The client deployed all their Clusters with VM-based Worker Nodes across six different physical servers.

For the Production and DR environments, based on the selected T-Shirt sizing, the Required Compute Capacity is 26 vCPUs and 91 GB RAM. Using the formula provided before, the Total Compute Capacity should be the following:

$$vCPU_{total} = \frac{26}{1 - \frac{1}{6}} = 31.2$$

$$RAM_{total} = \frac{91}{1 - \frac{1}{6}} = 109.2$$

Each Worker Node should have 5.2 vCPUs ($31.2/6$) and 18.2 GB RAM ($109.2/6$). Rounded up, the numbers are the following:

- 6 vCPU
- 20 GB RAM.

For the UAT and Dev combined Cluster, the Required Compute Capacity is 44 vCPU and 182 GB RAM. Using the formula provided before, the Total Compute Capacity should be the following:

$$vCPU_{total} = \frac{44}{1 - \frac{1}{6}} = 52.8$$

$$RAM_{total} = \frac{182}{1 - \frac{1}{6}} = 218.4$$

Therefore, each Worker Node should have 8.8 vCPUs ($52.8/6$) and 36.4 GB RAM ($218.4/6$). However, since the actual usage of the UAT and Dev environments will be much lower than the Production environment, the client can reduce it. The following numbers are reasonable for the six Worker Nodes of the UAT+Dev Cluster:

- 8 vCPU
- 32 GB RAM.

As mentioned in the *Finastra Kondor T-Shirt Sizing* section, Worker Nodes must have 100 GB of storage to hold the locally cached container images to accommodate deployment and upgrades of Finastra Kondor.

The Complete Picture

Based on the above decisions and calculations, this hypothetical client needs the following Production environment:

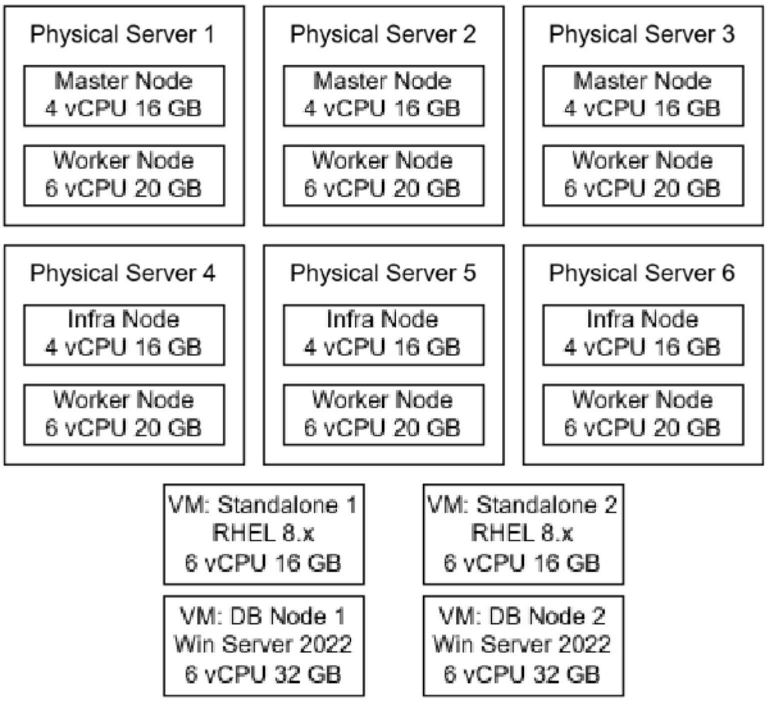

Figure 2: Production Cluster plus ancillaries

To cover every aspect, the Organization's existing Database Server also features VMs and a Standalone VM (with HA) hosting the Finastra License Server and additional interfaces. This Standalone VM will be used as well for IT administration tasks. The Citrix environment used to serve DealManager to the end users is not shown.

The DR environment is identical to the Production environment. Due to regulatory requirements, the DR environment is also required to provide full redundancy:

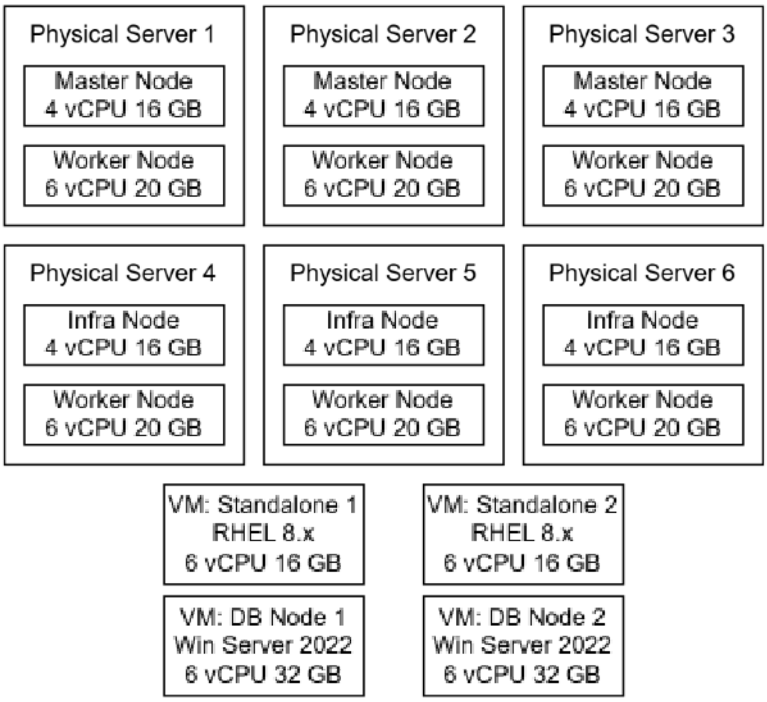

Figure 3: DR Environment

Finally, for the UAT + Dev environment, the example includes full redundancy for the UAT system too, so failover can be tested in a non-Production environment. Ultimately, it is the client’s decision whether this environment requires HA or not.

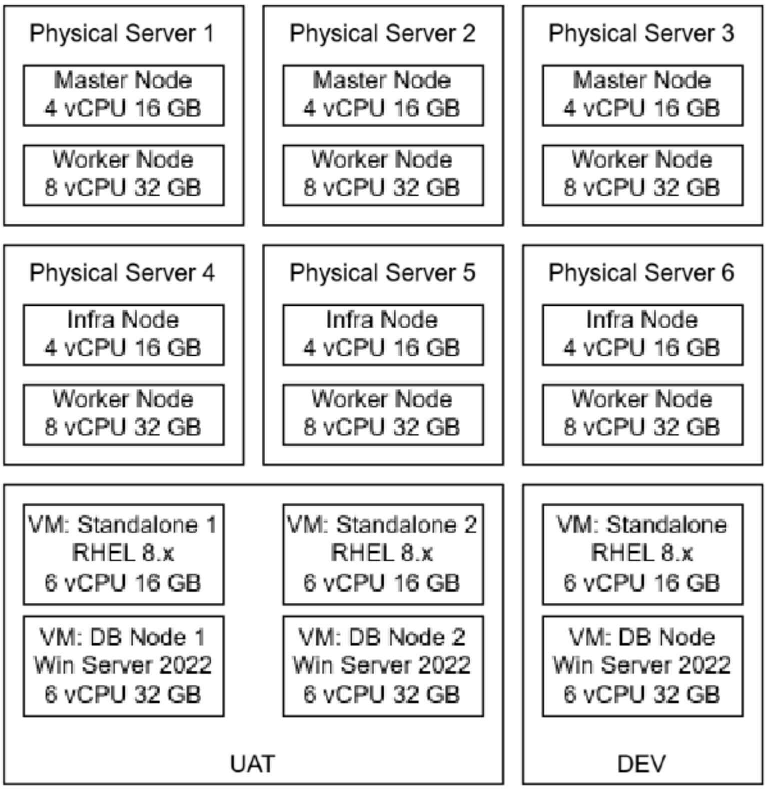

Figure 4: UAT + Dev Environment

Below is the total Hardware Bill of Materials:

	Number	Specifications
OpenShift Server	12	10 vCPU, 36 GB RAM 128 GB Storage for Master Nodes 512 GB Storage for Infra Nodes 256 GB Storage for Worker Nodes
OpenShift Server	6	12 vCPU, 48 GB RAM 128 GB Storage for Master Nodes 512 GB Storage for Infra Nodes 256 GB Storage for Worker Nodes
DB VM (existing)	7	6 vCPU, 32 GB RAM Windows Server 2022 Windows SQL Server 2022
Standalone VM (Can reuse from the existing infrastructure)	7	6 vCPU, 32 GB RAM Red Hat Enterprise Linux 8.x

Example 2

Client Overview

The Organization is an existing OpenShift user looking to migrate its Finastra Kondor to OpenShift.

The client is using Front Office, Back Office, Limits (Credit and Settlement Limits), Market Risk (Historical VaR, Stress Testing, Back Testing), and Market Limits System (MLS). They will not use Kondor Vision.

Their usage patterns are as follows:

Measure	Value
Concurrent Front Office Users	60
Concurrent Back Office Users	40
Concurrent Limit Users	25
Average Trades/Day	10,000 (mostly FX Spot)
Live Trades	7600
Number of Credit/Settlement Limits	<5000
Market Risk	
HS VaR	4 processes, 261 days, 4 risk classes, 250 days results archived
Stress Tests	5 Stress Test scenarios, 250 days results archived
Back Tests	4 Back Test scenarios, 250 days results archived

In addition, this Organization has the following existing infrastructure available:

- Existing OpenShift Platform Plus user
- Existing Microsoft SQL Server Database Servers using AlwaysOn

- Existing Rocket Software Exceed TurboX infrastructure for serving DealManager from a central location
- Existing DynaTrace deployment for metrics and log monitoring using the DynaTrace OneAgent operator
- Enterprise-Grade storage: Portworx
- F5 BIG-IP Load Balancer with F5 BIG-IP Container Ingress Services operator.

Since the Organization is using the OpenShift Platform Plus edition, they have access to Red Hat Quay enterprise-grade registry.

The client is planning to deploy Finastra Kondor inside *existing* Production, DR, and Non-Production clusters. Because of this, they will only need to extend the capacity of their Data Plane.

T-Shirt size selection

The Organization falls into the "Large" T-Shirt size category for Finastra Kondor F2B and the "Large" category for Risk. Therefore, they will need the following Kubernetes compute capacities:

Production & DR	**vCPU**	**RAM**
Kondor F2B	36	128
Risk – Standard Limits	9	28
Risk – Market Risk	6	16
Total	**51**	**172**

Note: As mentioned in the *Finastra Kondor T-Shirt Sizing* section, the sizing for Market Risk requires additional compute and memory resources.

The sizing above applies to the Production and DR environments. As for the UAT and Dev environments, a "Small" T-Shirt size can be the right choice.

UAT & DEV	**vCPU**	**RAM**
Kondor F2B	16	64
Risk – Standard Limits	6	19
Risk – Market Risk	4	16
Total	**26**	**99**

Note: As mentioned in the *Finastra Kondor T-Shirt Sizing* section, the sizing for Market Risk requires additional compute and memory resources.

Sizing of the Data Plane

The client has decided to deploy Finastra Kondor inside their *existing* Production, DR, and Non-Production clusters. These clusters currently have the following Data Plane configurations:

	Nodes	**vCPU**	**RAM**
Production Cluster	12	16	64
DR Cluster	12	16	64
Non-Production Cluster	12	16	64

This section proposes one method to calculate the additional resources required for extending the Data Plane to accommodate Finastra Kondor. Other approaches to this calculation are possible, too. If you are an OpenShift user, we recommend reaching out to your Red Hat Customer Advocate for further assistance in this Data Plane extension exercise.

The clusters below are sized to allow for a single Worker Node to fail while maintaining the required capacity (assuming an N + 1 model where the "+ 1" accounts for additional redundancy):

Required Capacity Existing Workload (excluding redundancy)	Required vCPU	Required RAM
Production Cluster	$(12 - 1) \times 16 = 176$	$(12 - 1) \times 64 = 704$
DR Cluster	$(12 - 1) \times 16 = 176$	$(12 - 1) \times 64 = 704$
Non-Production Cluster	$(12 - 1) \times 16 = 176$	$(12 - 1) \times 64 = 704$

The hypothetical Organization decides to extend these clusters to accommodate Finastra Kondor by adding extra Worker Nodes identical in configuration to the existing ones.

Below are the formulas to calculate the total number of Worker Nodes required to handle the existing workload, **plus** the Finastra Kondor workload.

Required Capacity Total Workload (excluding redundancy)	Required vCPU	Required RAM
Production Cluster	$176 + 51 = 227$	$704 + 172 = 876$
DR Cluster	$176 + 51 = 227$	$704 + 172 = 876$
Non-Production Cluster	$176 + 26 + 26 = 228$	$704 + 99 + 99 = 902$

Given the sizing of the existing and the additional Worker Nodes, the following additional number of Worker Nodes, *not counting the redundancy,* is needed for each environment:

Required Number of Worker Nodes (excluding redundancy)	Based on vCPU	Based on RAM
Production Cluster	$227 / 16 = 14.1$	$876 / 64 = 13.7$
DR Cluster	$227 / 16 = 14.1$	$876 / 64 = 13.7$
Non-Production Cluster	$228 / 16 = 14.2$	$902 / 64 = 14.1$

This means the following:

- The Production Cluster has **three** additional Worker Nodes - maximum of (14.1, 13.7) – 11 = 3.1
- The DR Cluster has **three** additional Worker Nodes – maximum of (14.1, 13.7) – 11 = 3.1
- The Non-Production Cluster has **three** additional Worker Nodes – maximum of (14.2, 14.1) – 11 = 3.2.

Note: To account for the additional redundancy requirements introduced by Finastra Kondor, the client needs to add four Worker Nodes to each Cluster.

As mentioned in the *Finastra Kondor T-Shirt Sizing* section, Worker Nodes must have 100 GB of storage to hold the locally cached container images to accommodate deployment and upgrades of Finastra Kondor.

The Complete Picture

Based on the above decisions and calculations, this client will need to scale their existing Production Cluster as follows:

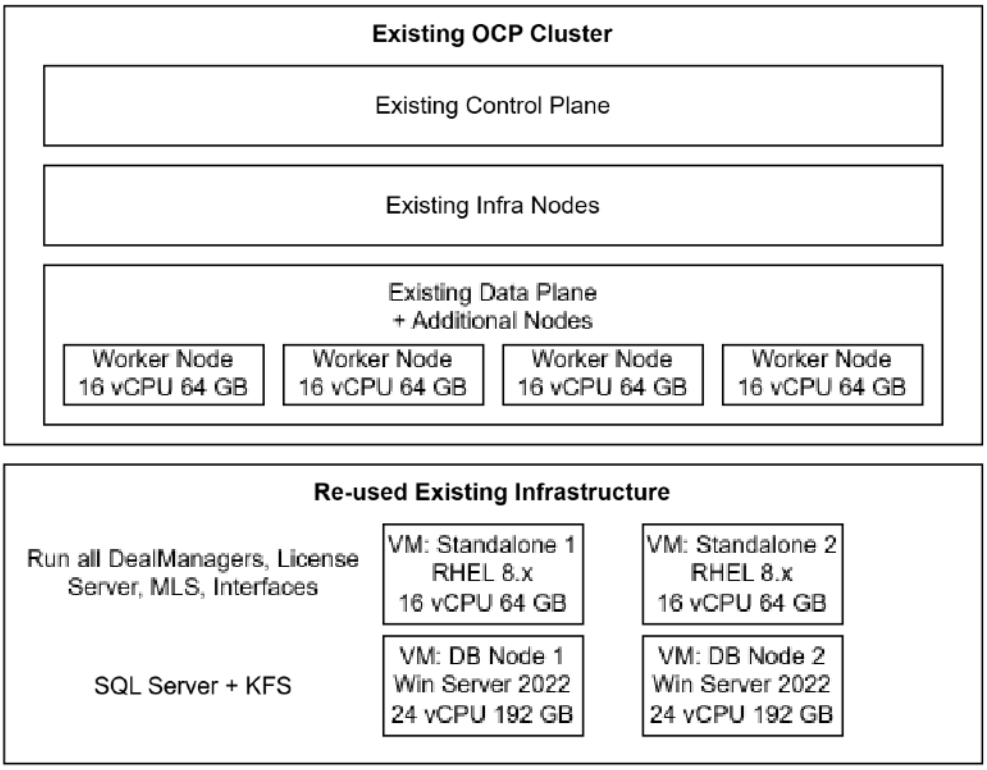
Existing OCP Cluster
Existing Control Plane
Existing Infra Nodes
Existing Data Plane + Additional Nodes
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Re-used Existing Infrastructure
Run all DealManagers, License Server, MLS, Interfaces
VM: Standalone 1 RHEL 8.x 16 vCPU 64 GB
VM: Standalone 2 RHEL 8.x 16 vCPU 64 GB
SQL Server + KFS
VM: DB Node 1 Win Server 2022 24 vCPU 192 GB
VM: DB Node 2 Win Server 2022 24 vCPU 192 GB

Figure 5: Production Cluster plus ancillaries

To cover every aspect, the Organization's existing Database Server also features VMs and a Standalone VM (with HA) hosting the Finastra License Server and additional interfaces. This Standalone VM will be used as well to run the DealManager instances that will be served to the End Users via Rocket Software Exceed TurboX.

The DR environment is identical to the Production environment. Due to regulatory requirements, the DR environment is also required to provide full redundancy:

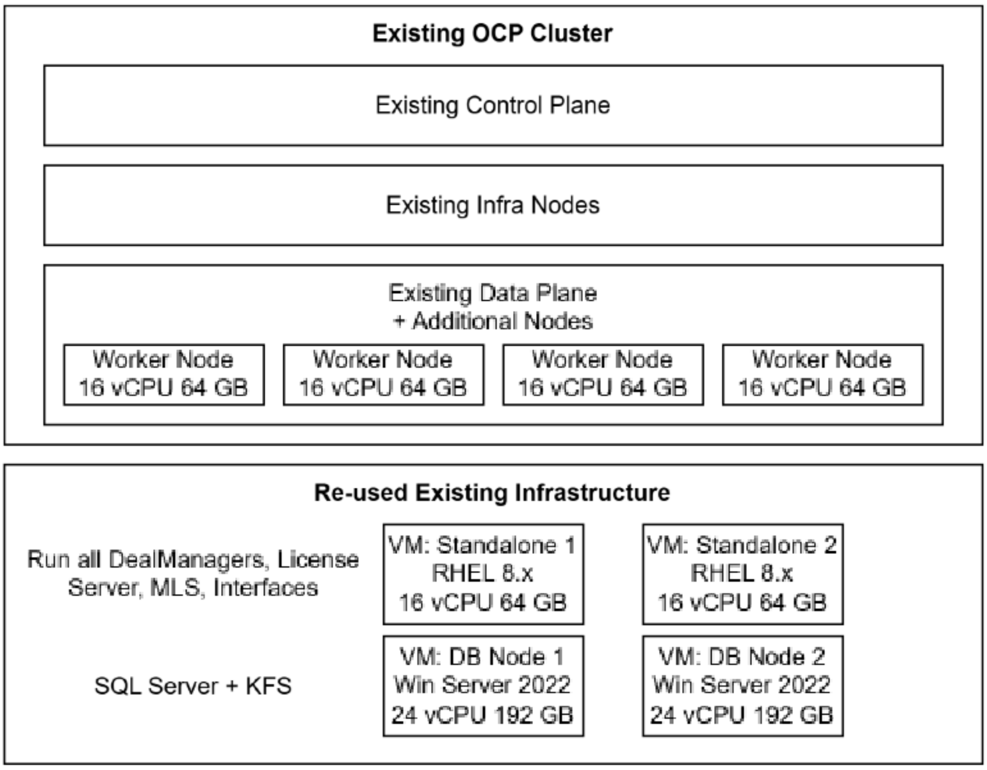
Existing OCP Cluster
Existing Control Plane
Existing Infra Nodes
Existing Data Plane + Additional Nodes
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Re-used Existing Infrastructure
Run all DealManagers, License Server, MLS, Interfaces
VM: Standalone 1 RHEL 8.x 16 vCPU 64 GB
VM: Standalone 2 RHEL 8.x 16 vCPU 64 GB
SQL Server + KFS
VM: DB Node 1 Win Server 2022 24 vCPU 192 GB
VM: DB Node 2 Win Server 2022 24 vCPU 192 GB

Figure 6: DR Cluster plus ancillaries

Note: To cover every aspect, the existing infrastructure for Microsoft SQL Server and Standalone VMs was added as well.

At the client's request, the UAT environment *does not* include High Availability.

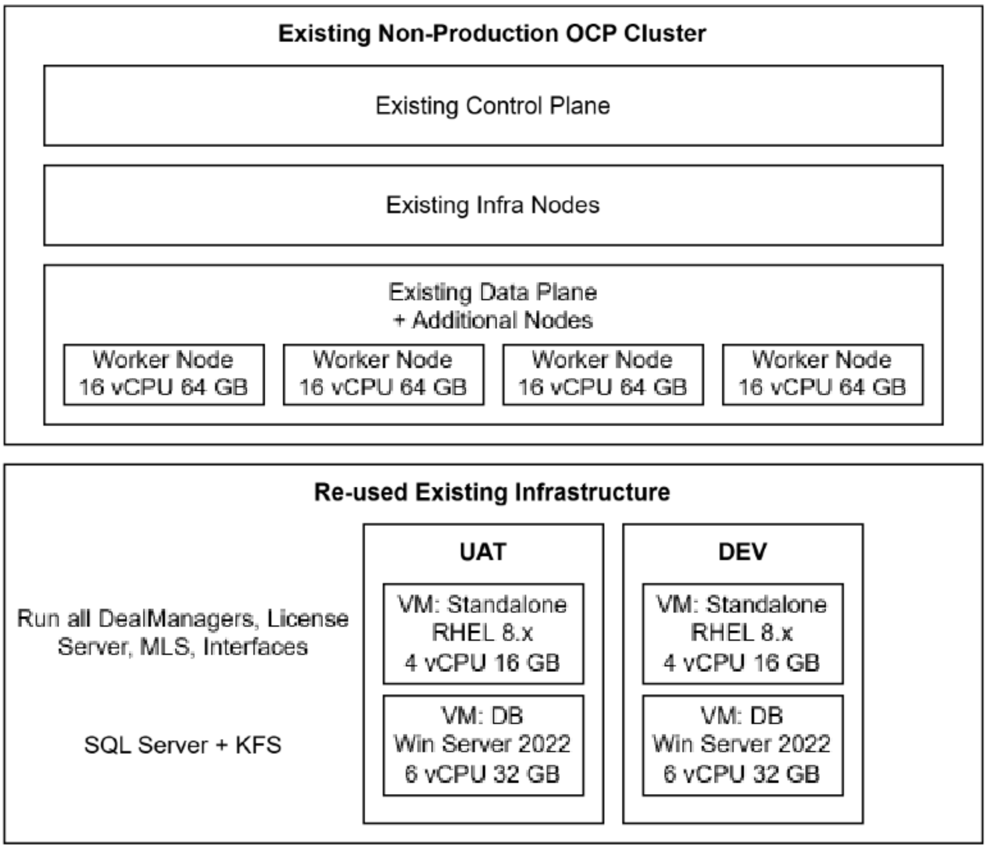
Existing Non-Production OCP Cluster
Existing Control Plane
Existing Infra Nodes
Existing Data Plane
+ Additional Nodes
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Worker Node 16 vCPU 64 GB
Re-used Existing Infrastructure
UAT
DEV
Run all DealManagers, License Server, MLS, Interfaces
SQL Server + KFS
VM: Standalone RHEL 8.x 4 vCPU 16 GB
VM: Standalone RHEL 8.x 4 vCPU 16 GB
VM: DB Win Server 2022 6 vCPU 32 GB
VM: DB Win Server 2022 6 vCPU 32 GB

Figure 7: UAT + Dev Cluster plus ancillaries

Below is the total Hardware Bill of Materials:

	Number	Specifications
OpenShift Worker Nodes (additional)	12	16 vCPU, 64 GB RAM 256 GB Storage
Prod & DR DB VM (existing)	4	24 vCPU, 192 GB RAM Windows Server 2022 Windows SQL Server 2022
UAT & Dev DB VM (existing)	2	6 vCPU, 32 GB RAM Windows Server 2022 Windows SQL Server 2022
Prod & DR Standalone VM (existing)	4	16 vCPU, 64 GB RAM Red Hat Enterprise Linux 8.x
UAT & Dev Standalone VM (existing)	2	4 vCPU, 16 GB RAM Red Hat Enterprise Linux 8.x

The future of digital finance can start here.
Consider our planet before printing.

About Finastra
Finastra is a global provider of financial software applications and marketplaces, and launched the leading open platform for innovation, FusionFabric.cloud, in 2017. It serves institutions of all sizes, providing award-winning software solutions and services across Lending, Payments, Treasury & Capital Markets and Universal Banking (Retail, Digital and Commercial Banking) for banks to support direct banking relationships and grow through indirect channels, such as embedded finance and Banking as a Service. Its pioneering approach and commitment to open finance and collaboration is why it is trusted by over 8,000 institutions, including 45 of the world's top 50 banks. For more information, **finastra.com**

Corporate Headquarters
4 Kingdom Street
Paddington
London W2 6BD
United Kingdom

T: +44 20 3320 5000