

Red Hat OpenShift and how to survive disaster of my data?

Marcin Steć
IBM

Zbigniew Parys
Red Hat





What is an app in container?

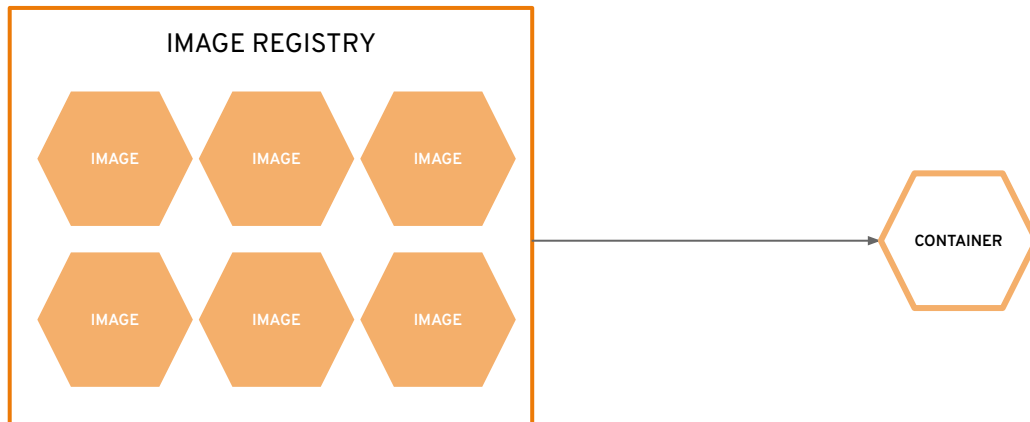
a container is the smallest compute unit



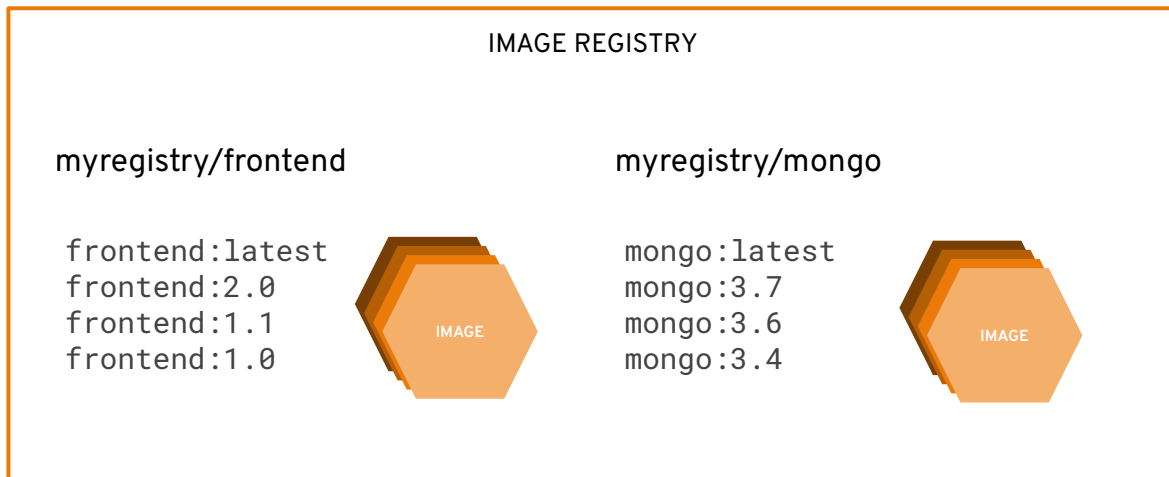
containers are created from container images




container images are stored in an image registry




an image repository contains all versions of an image in the image registry



containers are wrapped in pods which are units of deployment and management




10.140.4.44




10.15.6.55

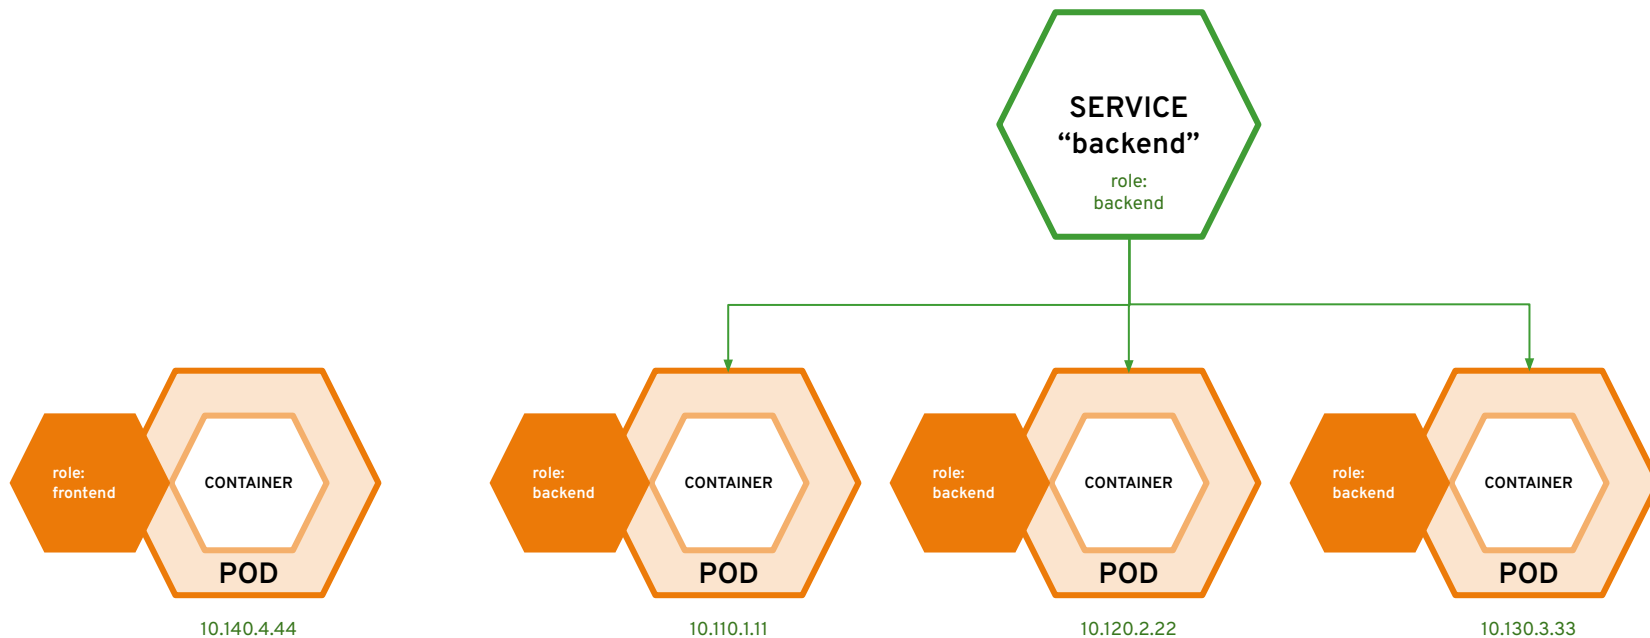
ReplicationControllers & ReplicaSets ensure a specified number of pods are running at any given time



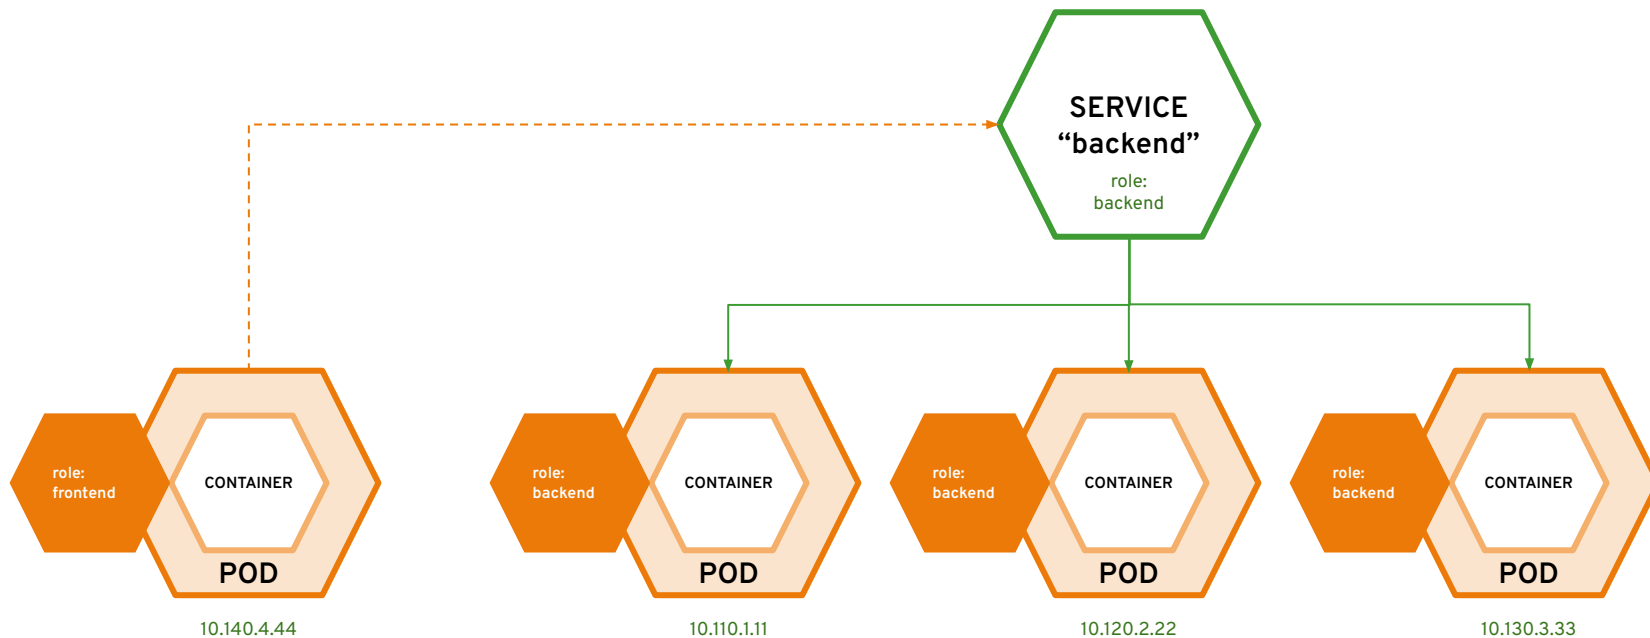
Deployments and DeploymentConfigurations define how to roll out new versions of Pods



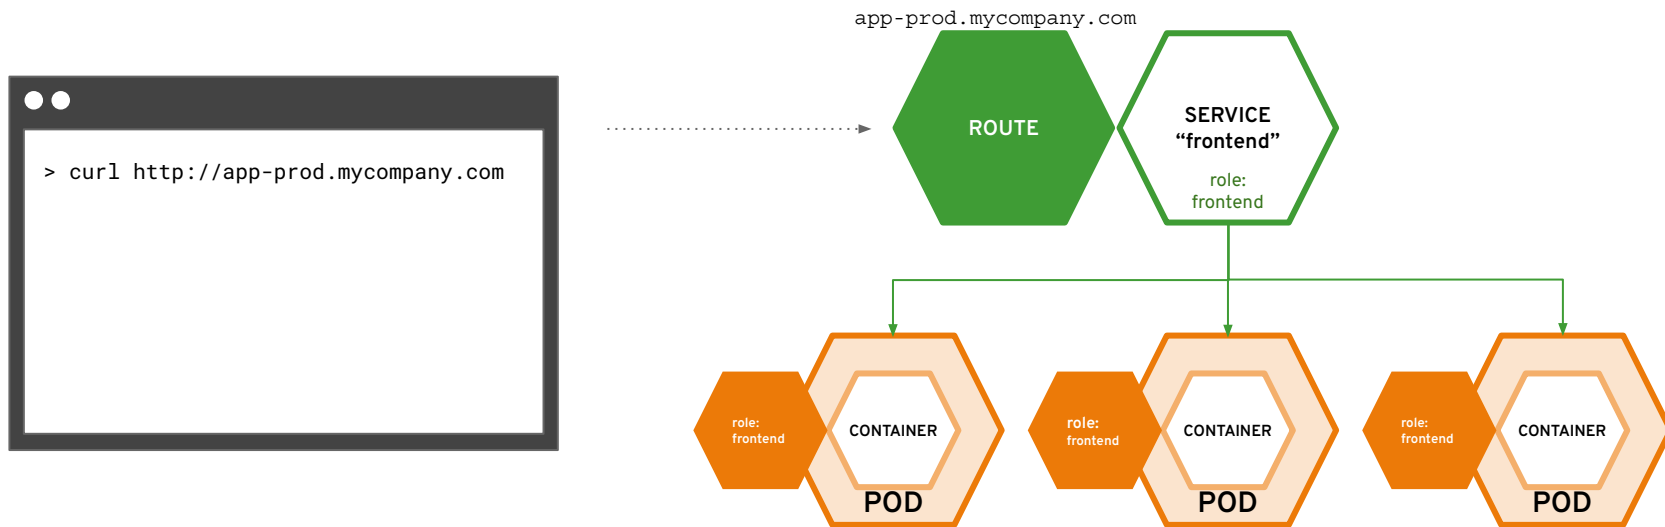
services provide internal load-balancing and service discovery across pods



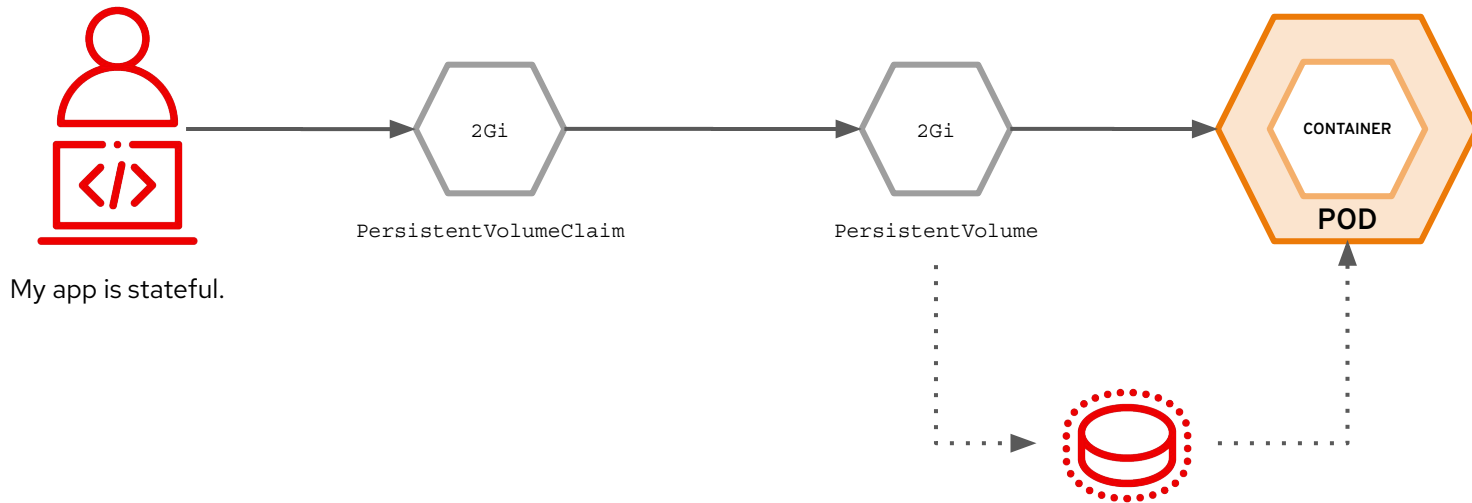
apps can talk to each other via services



routes make services accessible to clients outside the environment via real-world urls




Persistent Volume and Claims





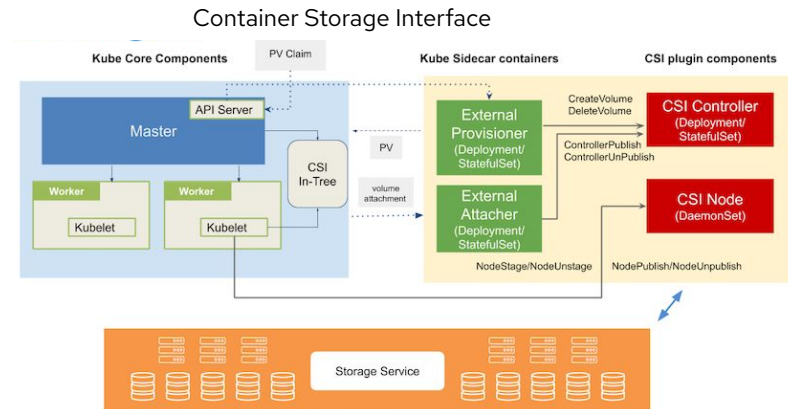
Kubernetes native storage & certified CSI plugins



Persistent Storage

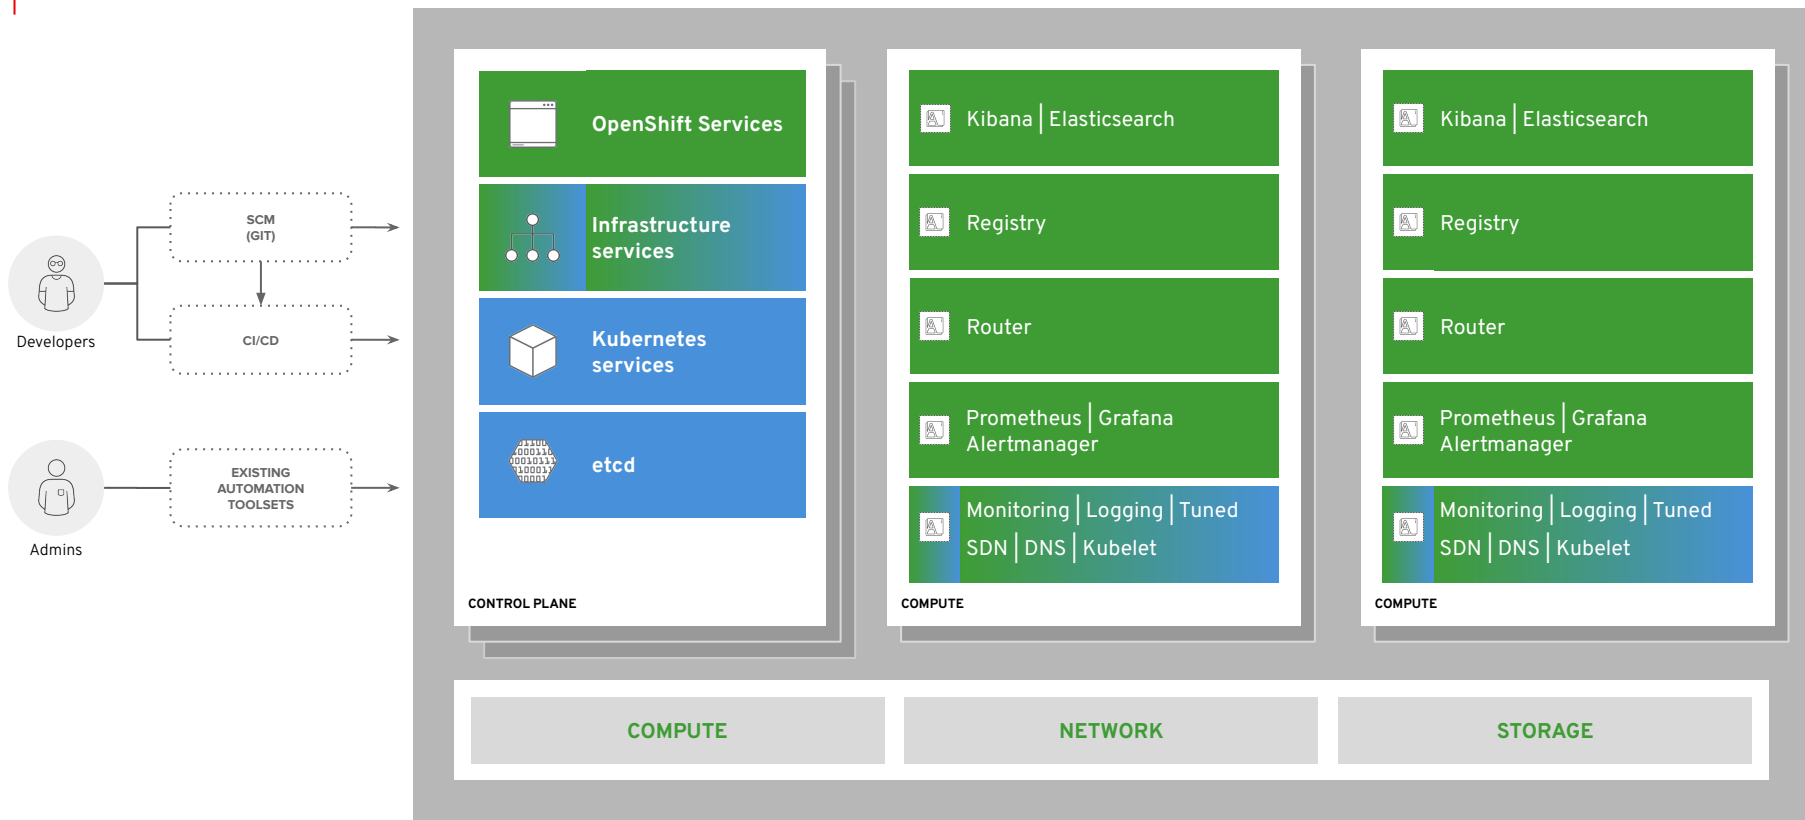
- RWO - ReadWriteOnce
- ROX - ReadOnlyMany
- RWX - ReadWriteMany

	RWO	ROX	RWX
AWS Elastic Block	X		
Azure Disk	X		
Azure File	X	X	X
IBM Spectrum Virtualize	X	X	
Fibre Channel	X	X	
GCE Persistent	X		
hostPath	X		
iSCSI	X	X	
Local volumes	X		
NFS	X	X	X
Red Hat OpenShift Data Foundation	X	X	X
VmWare vSphere	X		
IBM Spectrum Scale	X	X	X



CSI driver	CSI volume snapshots	CSI cloning	CSI resize
AWS EBS (Tech Preview)	✓	-	✓
Google Cloud Platform (GCP) persistent disk (PD) (Tech Preview)	✓	-	✓
OpenStack Cinder	✓	✓	✓
OpenStack Manila	✓	✓	✓
Red Hat Virtualization (oVirt)	-	-	-

we test and support the in-tree and select CSI provisioners, partners fill in the gaps for their storage, and we'll be (hopefully) seamlessly transitioning customers to CSI as in-tree drivers are removed.



App Definition and Development

Orchestration & Management

Runtime

Provisioning

Cloud

Database

Streaming & Messaging

Application Definition & Image Build

Continuous Integration & Delivery

Scheduling & Orchestration

Coordination & Service Discovery

Remote Procedure Call

Service Proxy

API Gateway

Service Mesh

Cloud-Native Storage

Container Runtime

Cloud-Native Network

Automation & Configuration

Container Registry

Security & Compliance

Key Management

Public

This landscape is intended as a map through the previously uncharted terrain of cloud native technologies. There are many routes to deploying a cloud native application, with CNCF Projects representing a particularly well-traveled path.

CLOUD NATIVE COMPUTING FOUNDATION

CLOUD NATIVE Landscape

Redpoint Amplify

Special

Platform

Certified Kubernetes - Distribution

Certified Kubernetes - Hosted

Certified Kubernetes - Installer

PaaS/Container Service

Observability and Analysis

Monitoring

Logging

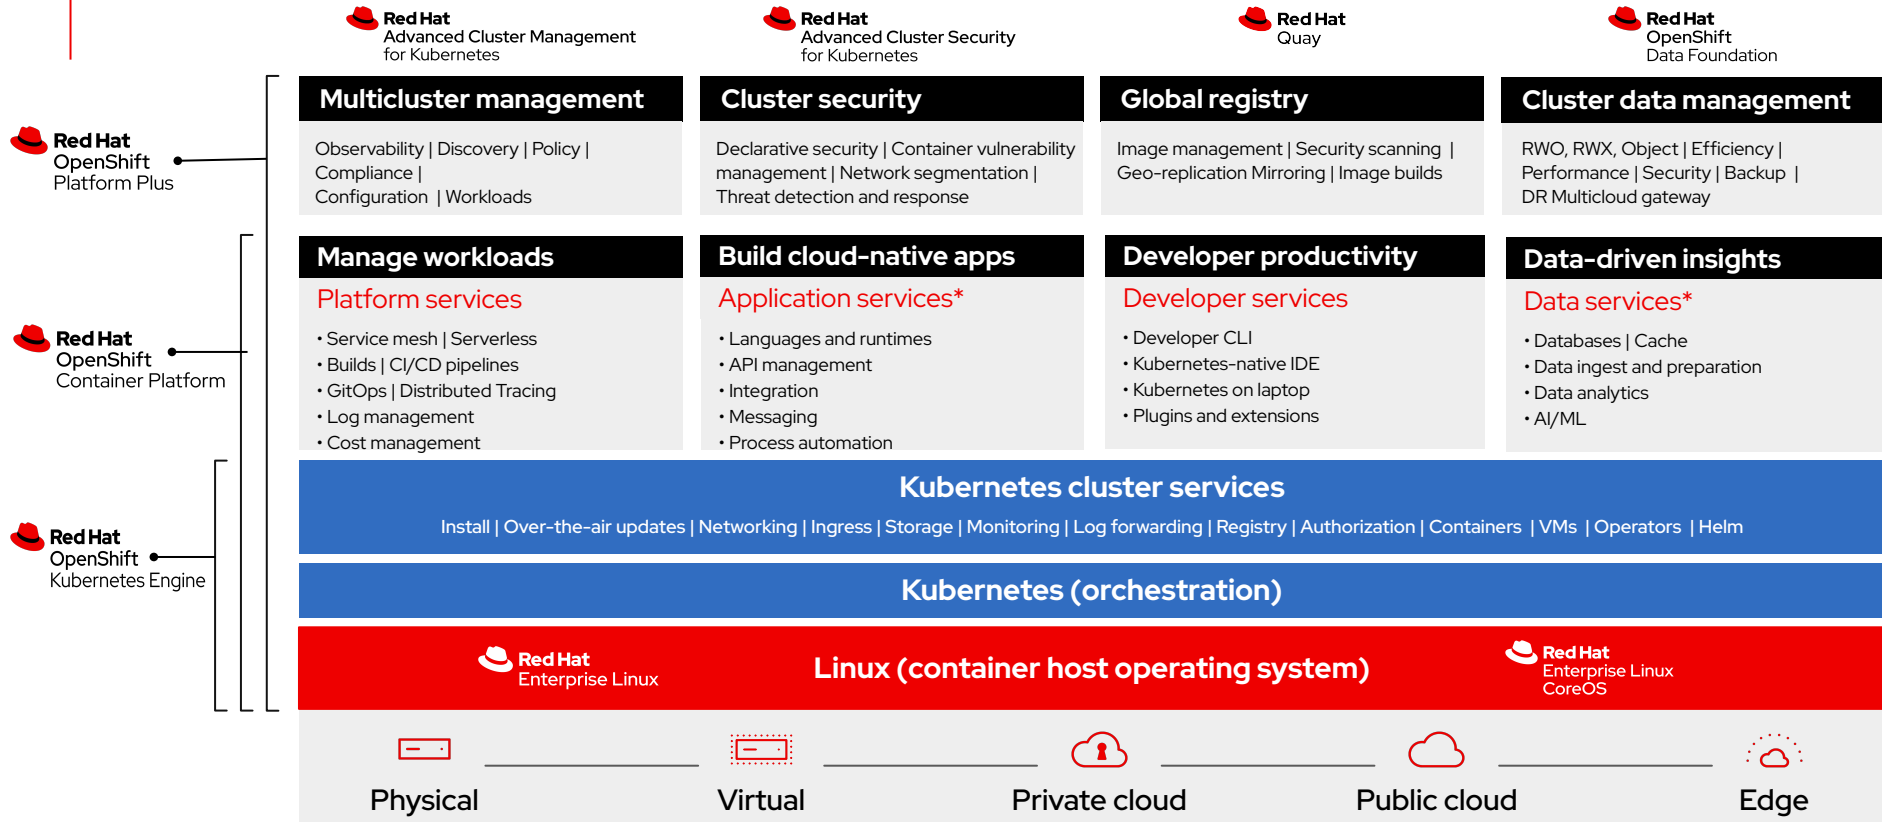
Tracing

Chaos Engineering

Serverless

Kubernetes Certified Service Provider

Kubernetes Training Partner

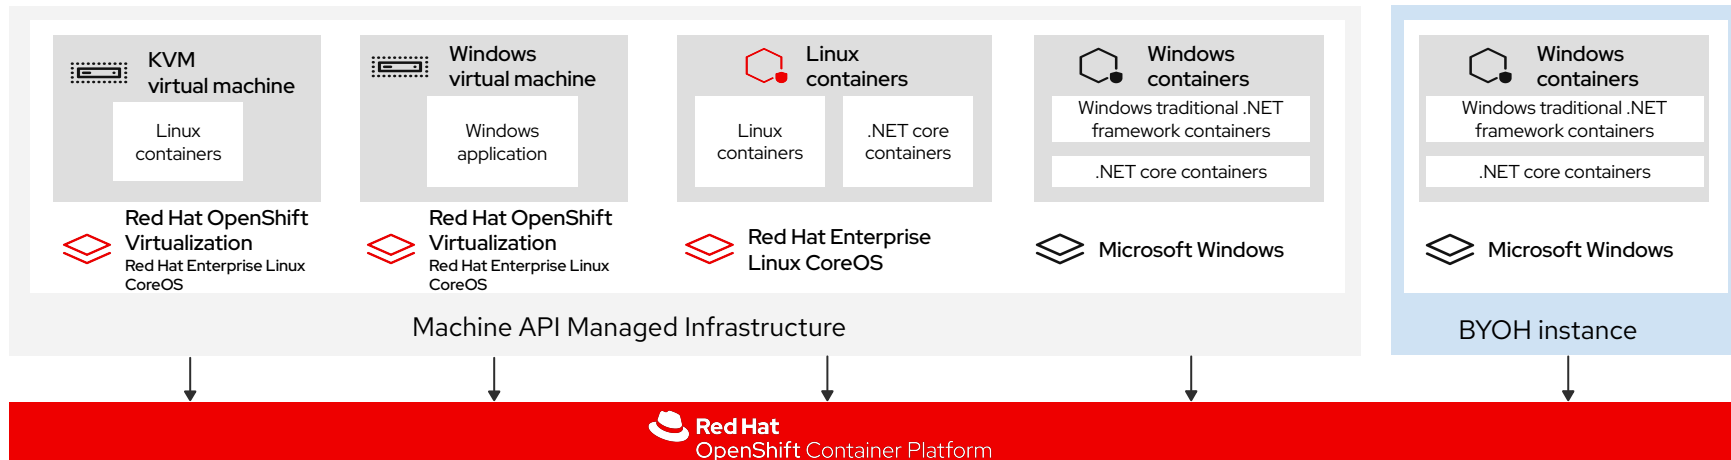


* Red Hat OpenShift® includes supported runtimes for popular languages/frameworks/databases. Additional capabilities listed are from the Red Hat Application Services and Red Hat Data Services portfolios.

** Disaster recovery, volume and multicloud encryption, key management service, and support for multiple clusters and off-cluster workloads requires OpenShift Data Foundation Advanced



Deploy and manage both Containers & VMs



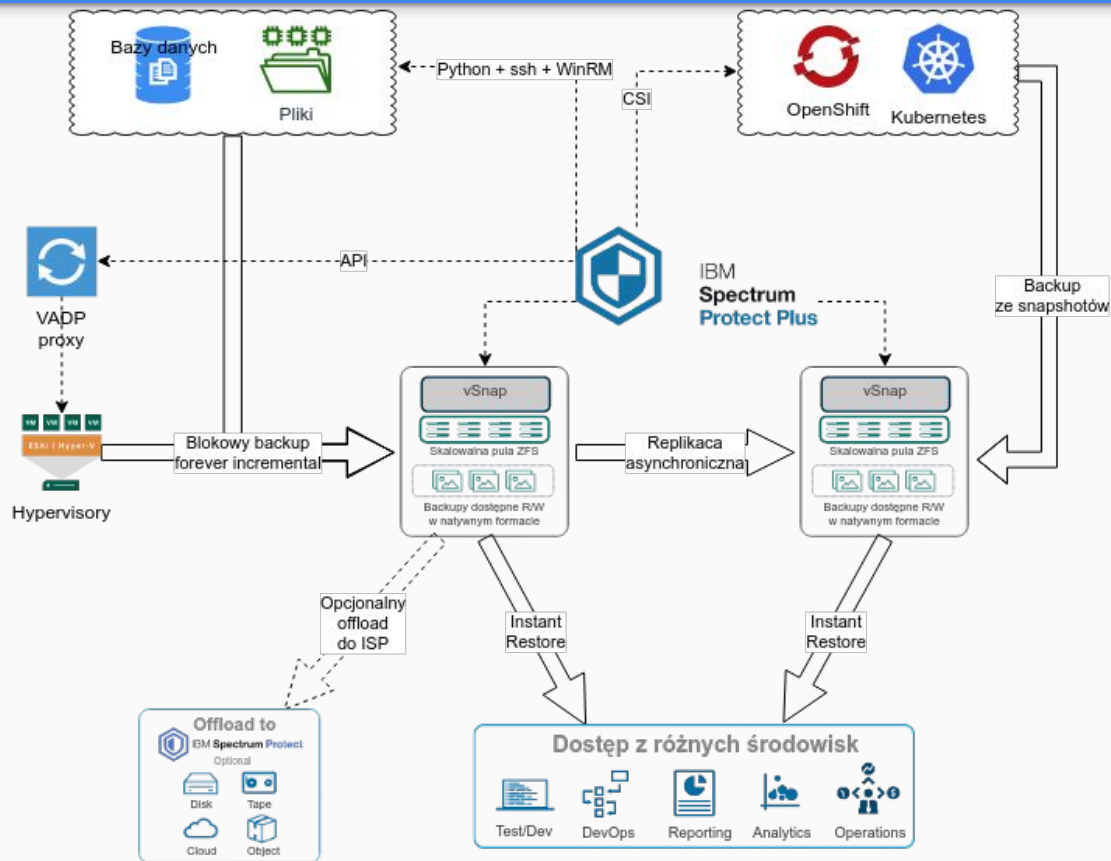
Red Hat OpenShift runs management and scheduling of:

- Linux containers on Red Hat Enterprise Linux CoreOS
- .NET core containers on Red Hat Enterprise Linux CoreOS or Windows
- Traditional **.NET framework containers on Windows**
- **Windows virtual machines with OpenShift Virtualization**

Bring your own Windows Hosts

Reuse "pet" Windows instances as OpenShift worker nodes, run Windows workloads and gain similar benefits that their Linux workloads get when being managed by OpenShift

Architektura SPP - przepływ danych



Integracja z IBM Spectrum Protect

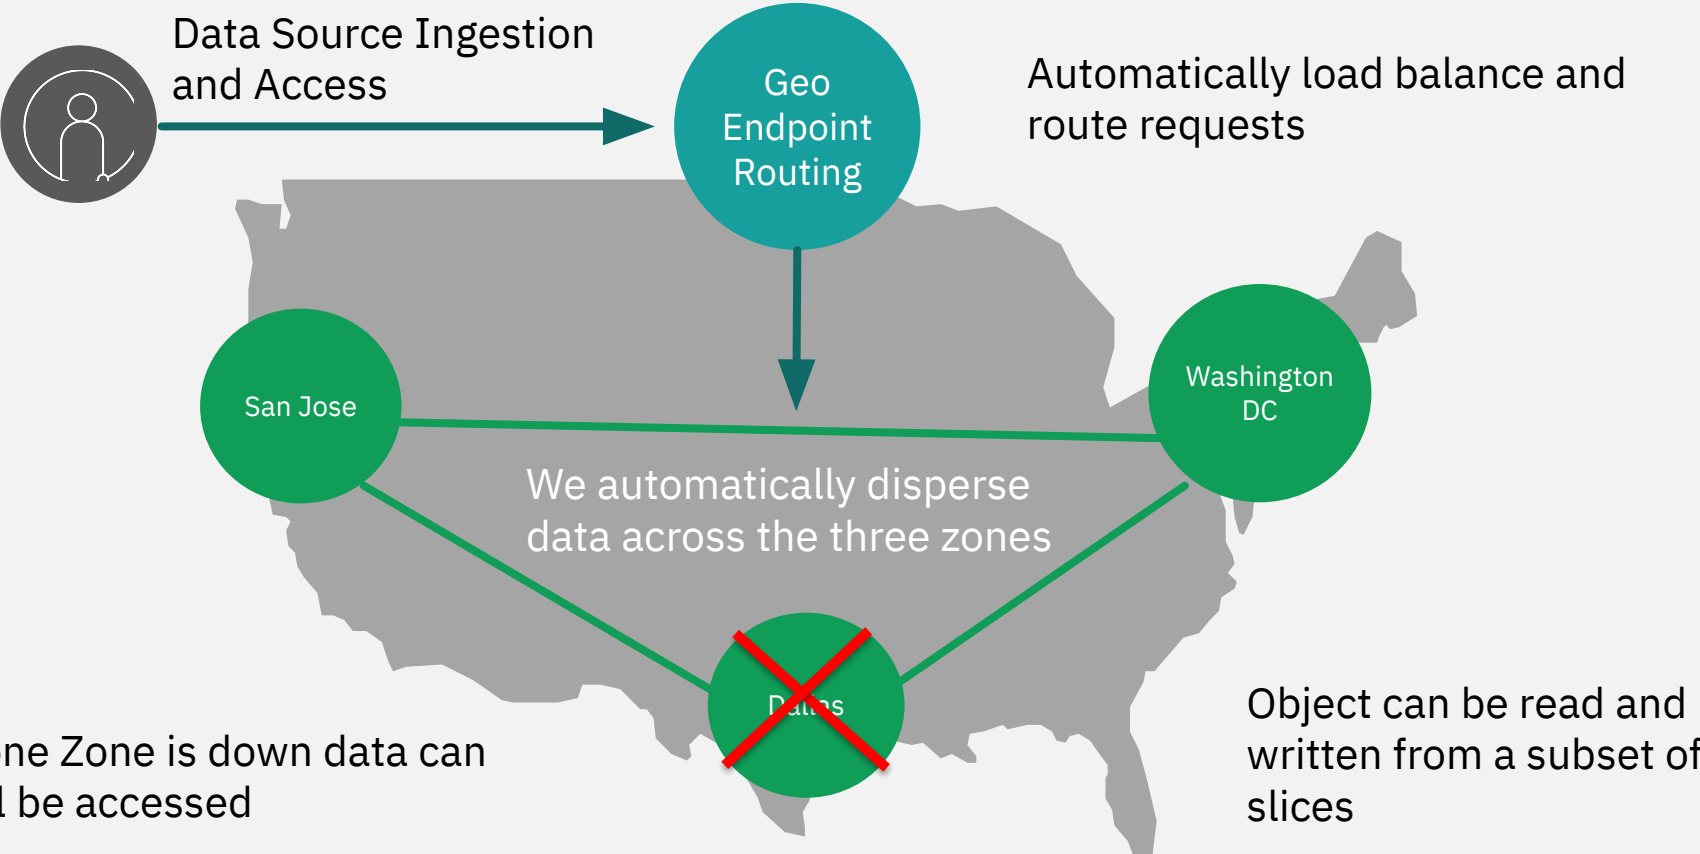
Spectrum Protect

- Na serwerze Protect musi być skonfigurowany i uruchomiony *Object Agent*
- Klienci obiektowi w specjalnych domenach
- Management klasy i copy grupy tworzone automatycznie
- Dane składowane w dwóch typach pul:
 - Standard - musi być pulą kontenerową (dir lub cloud)
 - Cold - Musi być pulą typu cold-data-cache (plikowa z *next* na taśmy lub VTL)

Spectrum Protect Plus

- Klient S3
- Dwa typy requestów: S3 i S3 Glacier
- Glacier to offload do puli *coldpool* czyli na taśmy.

Cross Region Resiliency



Storage Classes and Archive

For long-term data retention, rarely or never accessed

For any workload, ideal for unknown or changing data activity



Automatic classification and cost optimization based on data activity using built in tiers (hot, cool, cold)

For predictable and consistent data activity



For active data, accessed multiple times a month



For less active data, accessed once a month or less



For cold data, accessed a few times a year

For long-term data retention



For lowest cost with restore ≤ 12 hours

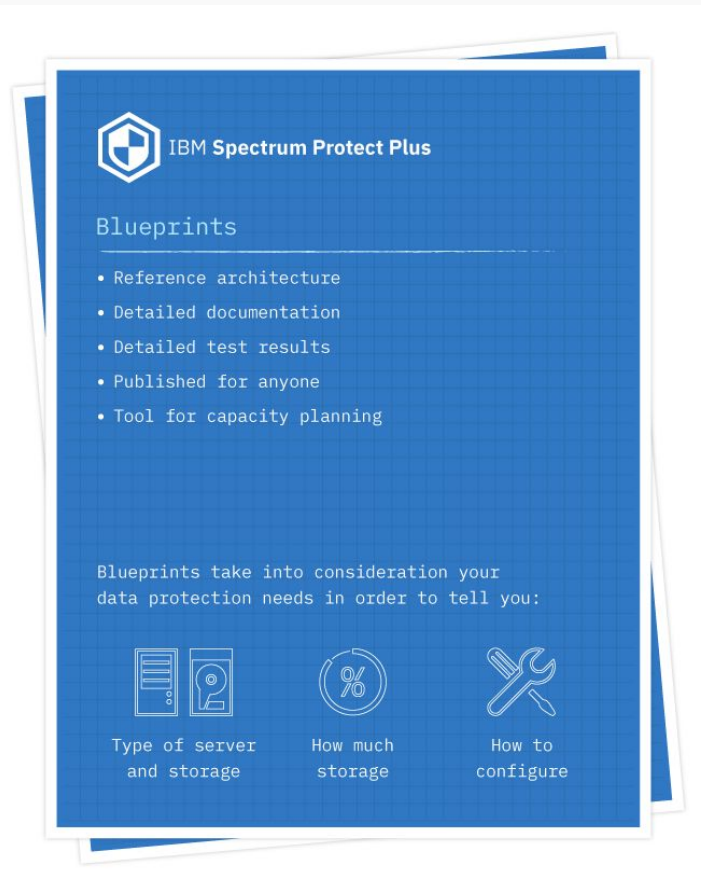


For faster access with restore ≤ 2 hours

IBM Blueprint: Wspomaganie projektowania

IBM Blueprint dla Spectrum Protect Plus

- Referencyjne architektury dla kilku rozmiarów rozwiązań.
- Podręcznik budowy uwzględniający wydajność i HA.
- Arkusz kalkulacyjny pozwalający wyliczyć wymagany sprzęt i oprogramowanie.



IBM Blueprint: Wspomaganie projektowania - arkusz wymiarujący



IBM Spectrum Protect Plus

Global options

vSnap enable compression	Yes
Compression estimate ratio x:1	2 :1
vSnap enable deduplication	No
Deduplication estimate ratio x:1	1,5 :1
Enable vSnap encryption	No

VADP transport method	NBDSSL
-----------------------	--------

VMware host to VADP proxy network (Gbps)	10
vSnap network (Gbps)	10


Optional replication input workload

Peak replication rate (Gbps)	0
Replication physical capacity (TiB)	0
Replication annual growth rate %	0,0%

Sizing estimates are based off of Spectrum Protect Plus Version 10.1.6

How to: Begin with the global variables at the "Start Here" tab, then select the applicable client tabs and fill in the "front end" values at minimum. View summarized results on the "Sizing Results" tab.

Total data reduction estimate	50%
-------------------------------	-----



Click to reset

Założenia wstępne:

- Kompresja
- Poziom deduplikacji
- Szyfrowanie
- Transport danych
- Szczytowe wymagania na pasmo
- Prognozowane wzrosty danych

IBM Blueprint: Wspomaganie projektowania - arkusz wymiarujący

Zakładki z klientami:

- Dla każdego typu klienta cztery polityki
- VMware i Hyper-V predefiniowane
- Możliwość rekonfiguracji

VMware

Daily change rate: 5.0%
Annual growth: 10.0%

At minimum, change the values in the orange boxes

Policy	Gold	Silver	Bronze	Custom
Front-end (TiB)	318,00	203,00	0,00	0,00
Backup frequency	1	1	1	1
Day(s)	Day(s)	Day(s)	Day(s)	Week(s)
Backup duration Hr(s)	6	6	8	8
Retention day(s)	35	13	35	35
Replication policy?	Yes	Yes	No	No
Replication frequency	1	1	1	1
Day(s)	Day(s)	Day(s)	Day(s)	Week(s)
Daily replication duration hr(s)	3	3	3	3
Replication retention day(s)	35	13	35	35
Copy to standard object storage policy?	Yes - Primary vSnap	No	No	No
Copy to object storage frequency	1	1	1	1
Day(s)	Day(s)	Day(s)	Day(s)	Week(s)
Copy to object storage duration hr(s)	3	3	3	3
Copy to object storage retention day(s)	3650	3650	3650	3650


Primary Copy								
	1 year projection	2 year projection	3 year projection	4 year projection	5 year projection	6 year projection	7 year projection	8 year projection
Peak vSnap front-end rate (MiB/s)	1265	1391	1530	1683	1852	2037	2240	2464
Peak vSnap back-end rate (MiB/s)	632	696	765	842	926	1018	1120	1232
Peak copy to object stg. rate (MiB/s)	772	849	934	1027	1130	1243	1367	1504
Front-end data (TiB)	521,0	573,1	630,4	693,5	762,8	839,1	923,0	1015,3
Managed data (TiB)	1183,4	1301,7	1431,9	1575,1	1732,6	1905,9	2096,5	2306,1
Compression savings %	50%	50%	50%	50%	50%	50%	50%	50%
Compression savings (TiB)	591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1
Dedup savings %	0%	0%	0%	0%	0%	0%	0%	0%
Dedup savings (TiB)	0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
Data reduction savings %	50%	50%	50%	50%	50%	50%	50%	50%
Data reduction savings (TiB)	591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1

Start Here | VMware | Hyper-V | Application-1 | Application-2 | Application-3 | Application-4 | Sizing Results | Version History

IBM Blueprint: Wspomaganie projektowania - arkusz wymiarujący

Podsumowanie wymiarowania:

- Ilość danych do zalicencjonowania
- Prognoza wzrostu na 8 lat.
- Liczba vSnapów do spełnienia wymagań wydajnościowych, pojemnościowych oraz dostępności (repliki)



Preferred vSnap size (TiB)	Automatic
vSnap reserve %	Automatic
VADP proxy resides on vSnap server?	Yes

Global Preferences - Protection recommended	
Target free space error (percentage)	10
Target free space warning (percentage)	20

Replication workload	
Peak replication rate (Gbps)	10,61
Replication physical capacity (TiB)	650,67
Replication annual growth rate %	10,0%

		1 year projection	2 year projection	3 year projection	4 year projection	5 year projection	6 year projection	7 year projection	8 year projection
Primary Copy									
Peak vSnap front-end rate	Gbps	10,6	11,7	12,8	14,1	15,5	17,1	18,8	20,7
Peak vSnap back-end rate	MiB/s	632,3	695,5	765,1	841,6	925,8	1018,3	1120,2	1232,2
Front-end data (TiB)		521,0	573,1	630,4	693,5	762,8	839,1	923,0	1015,3
Managed data (TiB)		1183,4	1301,7	1431,9	1575,1	1732,6	1905,9	2096,5	2306,1
Compression savings %		50%	50%	50%	50%	50%	50%	50%	50%
Compression savings (TiB)		591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1
Dedup savings %		0%	0%	0%	0%	0%	0%	0%	0%
Dedup savings (TiB)		0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
Data reduction savings %		50%	50%	50%	50%	50%	50%	50%	50%
Data reduction savings (TiB)		591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1
vSnap size estimate - data (TiB)		591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1
vSnap size estimate - logs (TiB)		0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
vSnap physical capacity estimate (TiB)		591,7	650,9	716,0	787,6	866,3	952,9	1048,2	1153,1
vSnap reserve %		10%	10%	10%	10%	10%	10%	10%	10%
vSnap size estimate - reserve (TiB)		59,2	65,1	71,6	78,8	86,6	95,3	104,8	115,3
vSnap base capacity		650,9	716,0	787,6	866,3	952,9	1048,2	1153,1	1268,4
Replication physical capacity estimate (TiB)		0,0	0,0	0,0	0,0	0,0	0,0	0,0	0,0
vSnap physical capacity estimate (TiB)		650,9	716,0	787,6	866,3	952,9	1048,2	1153,1	1268,4
vSnap quantity		4	4	4	5	5	6	6	7
per vSnap requirements									
peak vSnap front-end rate	Gbps	2,7	2,9	3,2	2,8	3,1	2,8	3,1	3,0
peak vSnap back-end rate	MiB/s	158,1	173,9	191,3	168,3	185,2	169,7	186,7	176,0
vSnap pool size (TiB)		200	200	200	200	200	200	200	200
vSnap CPU cores		18	18	18	18	18	18	18	18
vSnap memory (GB)		56	56	56	56	56	56	56	56
vSnap log SSD; optional (GiB)		10	10	10	10	10	10	10	10
vSnap cache SSD; optional (GiB)		500	500	500	500	500	500	500	500
vSnap cloud cache SSD (GiB)		1536	1536	1536	1536	1536	1536	1536	1536
per VADP requirements									
VADP quantity		0	0	0	0	0	0	0	0
VADP CPU cores		0	0	0	0	0	0	0	0
VADP memory (GB)		0	0	0	0	0	0	0	0
VADP Softcap		12	12	12	12	12	12	12	12
Replication Copy									

Instalacja

IBM Spectrum Protect Plus

Instalacja OVA w środowisku ESX (jest także H-V)

Rozruch vSnap.
Instalacja dodatkowych wg. potrzeb

Określenie polityk ochrony danych.
Można użyć wbudowanych

Dodanie hypervisorów i klientów Bare Metal



Inicjalne ustawienia, np hasła.

Utworzenie VADP proxy

OVA

1. Zaznacz klaster.
2. Kliknij *Actions*.
3. Wybierz *Deploy OVF Template*.

The screenshot displays the vSphere Client interface. The left-hand navigation pane shows a tree view of the environment. The cluster '5817_IIC' is selected and highlighted with a red circle labeled '1'. The main pane shows the 'Summary' tab for this cluster, with the 'ACTIONS' menu open, also highlighted with a red circle labeled '2'. Within the 'ACTIONS' menu, the 'Deploy OVF Template...' option is highlighted with a red circle labeled '3'. The interface also shows various system alerts at the bottom, such as 'Host memory' and 'Datastore: Datastore usage'.

1. Podaj detale konfiguraciji:
 - a. IP
 - b. Maska
 - c. Gateway
 - d. DNS
 - e. domena
2. Kliknij *Next*.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values

▼ Hostname Configuration 1 settings

Hostname

The Hostname or FQDN. The blank default name is localhost.

spp.ms6018.iic.pl.ibm.com

▼ Connection Configuration 5 settings

Network IP Address

The IP address for this interface. Leave blank if DHCP is desired.

10.10.13.18

Network Prefix

The prefix for this interface. Typical values are 8,16 or 24. Netmask in old formats like 255.255.255.0 are not supported. Check with your network administrator. Leave blank if DHCP is desired.

255.255.0.0

Default Gateway

CANCEL

BACK

NEXT

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 [linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

 [facebook.com/redhatinc](https://www.facebook.com/redhatinc)

 twitter.com/RedHat