

Welcome to: Brunch and Learn

Today's session will begin shortly

Friday 21st June 2024


Multi-Architecture Computing with OpenShift

Paul Chapman: Global Power Modernisation Technical Lead



Note:

Upon joining, you are muted and cannot see other attendees
Feel free to use the Chat or Q&A functions (panel on RHS of screen)



IBM UKI Brunch & Learn

Webinar

21 June 2024

Red Hat OpenShift Multi-Arch Compute


Deploy Microservice Solution



[Paul Chapman](#)

IBM, Global Power Modernization Technical Lead

With credit to Paul Bastide & Multi-Arch Dev Team



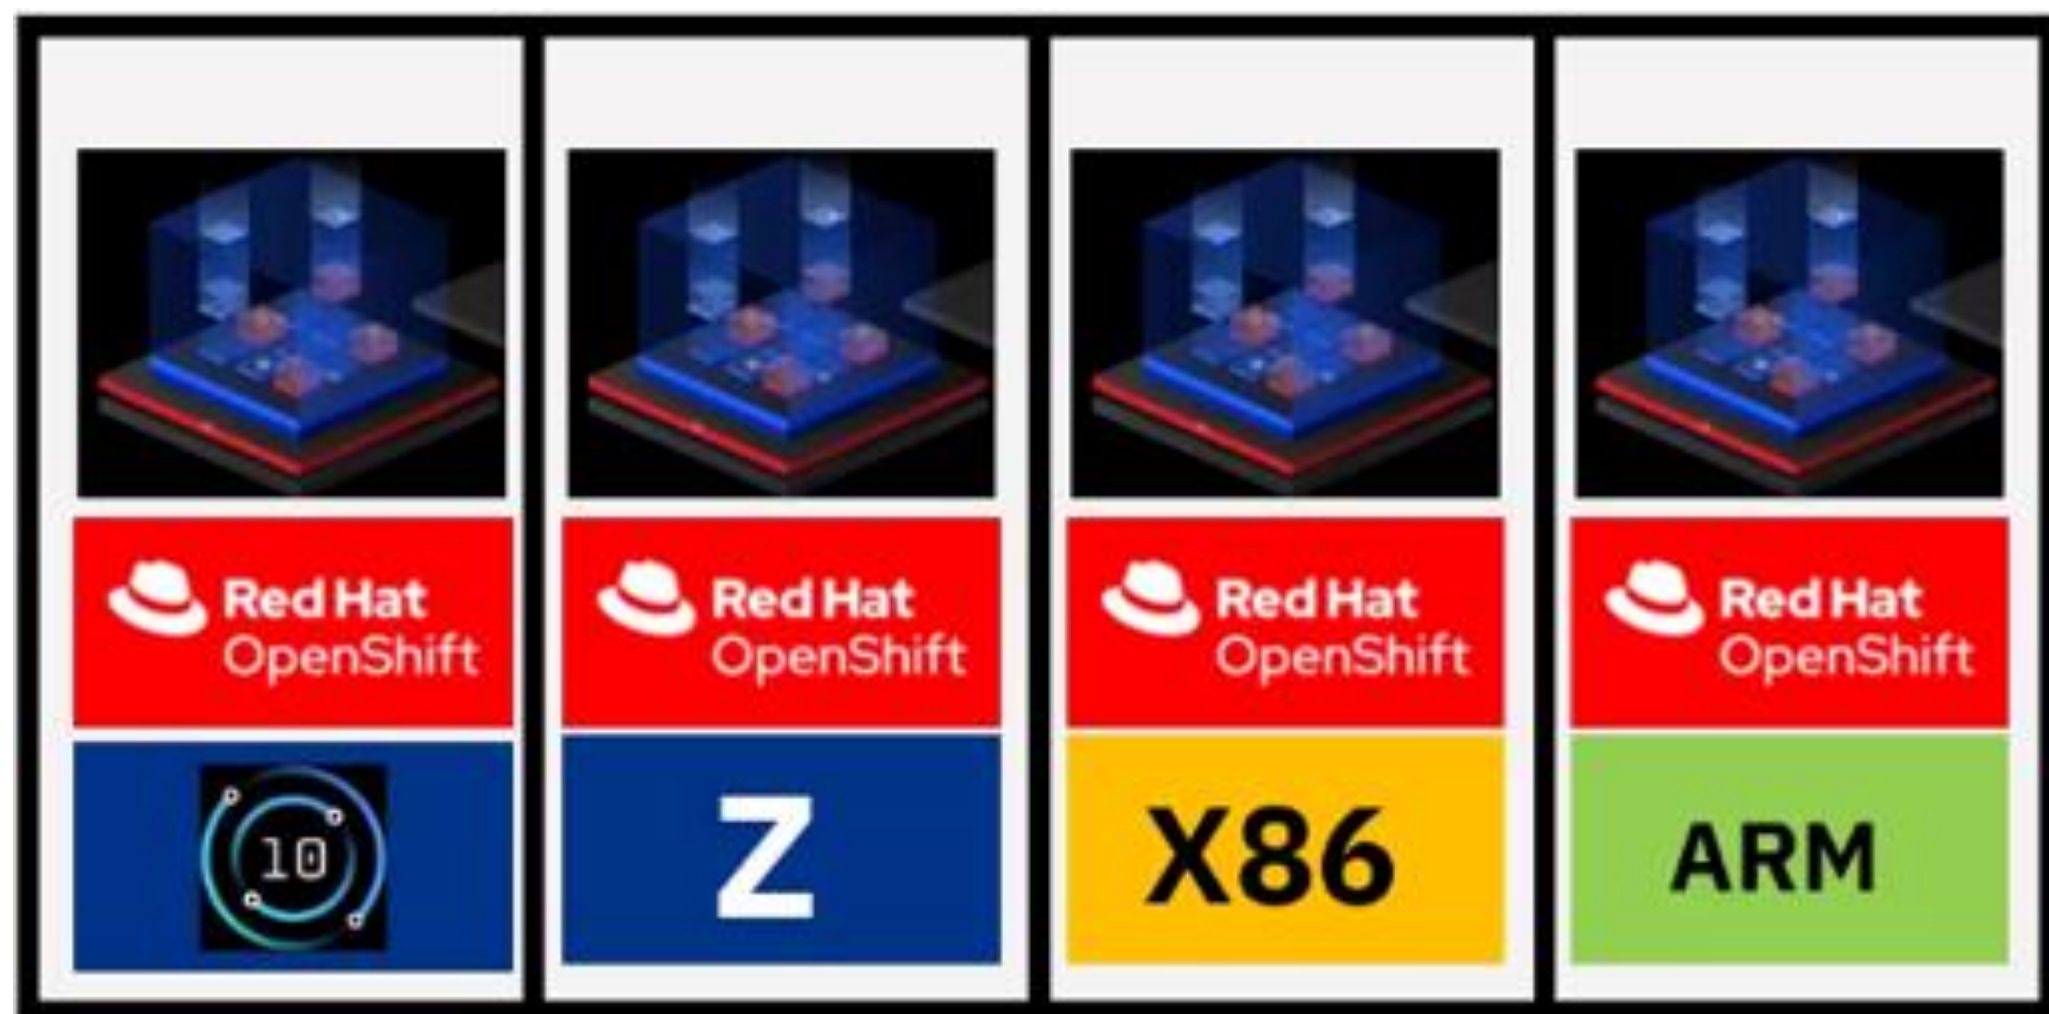
Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

Multi-Arch History

Previously

A cluster was required for each architecture




Which complicates:

1. Delivery of solutions requiring multi-arch
2. Migration to alternative architectures

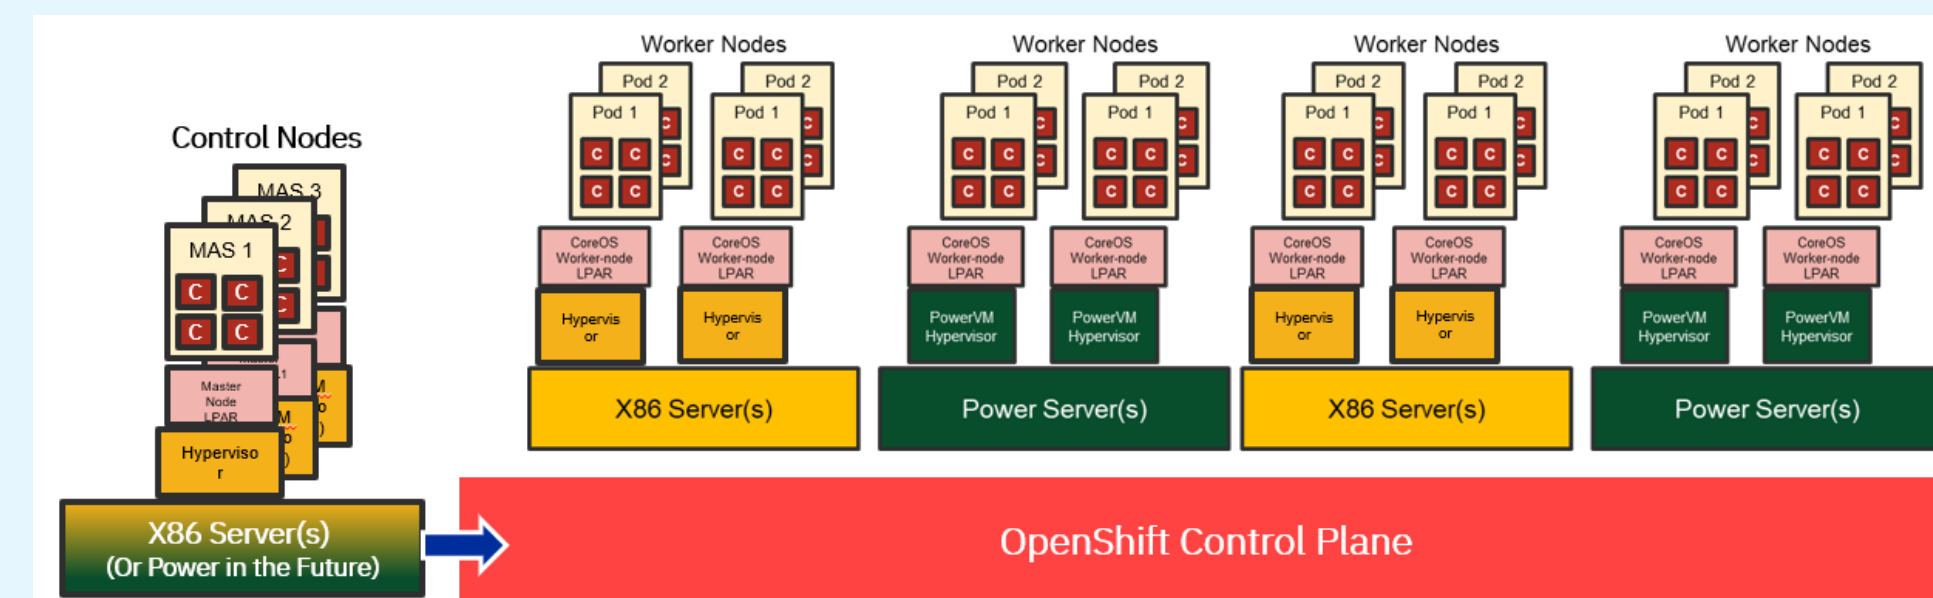
Now

Single cluster supports multiple architectures



4.14 Control Plane must reside on x86

4.15 Control Plane may reside on Power



Multi-Arch History

GA on Power Nov 2023

Multi-Arch available on Power with OCP 4.14



Paul Chapman • You
Global Power Modernisation Technical Lead
2w • Edited •

OpenShift Container Platform 4.14 is now available on IBM Power ppc64le, which enables:

Multi-Architecture Compute with Power

Single Node OpenShift

Hosted Control Plane – Technology Preview

oc-mirror plugin

IPI for PowerVS continues in Technology Preview

<https://lnkd.in/e3EwJgij>

[#powermod](#) [#ibmpowersystems](#) [#redhatopenshiftcontainerplatform](#) [#power](#)

OpenShift Container Platform 4.14 release notes

docs.openshift.com • 2 min read

OpenShift Container Platform (RHSA-2023:1326) is now available. This release uses Kub...

Andrea Corbelli and 83 others

2 comments • 16 reposts

https://www.linkedin.com/posts/chapmanp_openshift-container-platform-414-release-activity-7125400057436073984-68Ju?utm_source=share&utm_medium=member_desktop

Red Hat OCP 4.14 Release Notes

The screenshot shows the Red Hat OpenShift documentation page for the 4.14 release notes on IBM Power. The page features a navigation menu on the left with categories like 'About', 'Release notes', 'Getting started', etc. The main content area is titled 'IBM Power' and includes an 'IMPORTANT' warning box stating that compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS). Below this, there are sections for 'IBM Power notable enhancements' and a list of features introduced in this release, including 'Multi-architecture compute nodes' and 'oc-mirror plugin'.

https://docs.openshift.com/container-platform/4.14/release_notes/ocp-4-14-release-notes.html

Agenda

- 01 Multi-Arch Compute History
- 02 [IBM Power Strategy](#)
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

IBM & Power Strategy

1

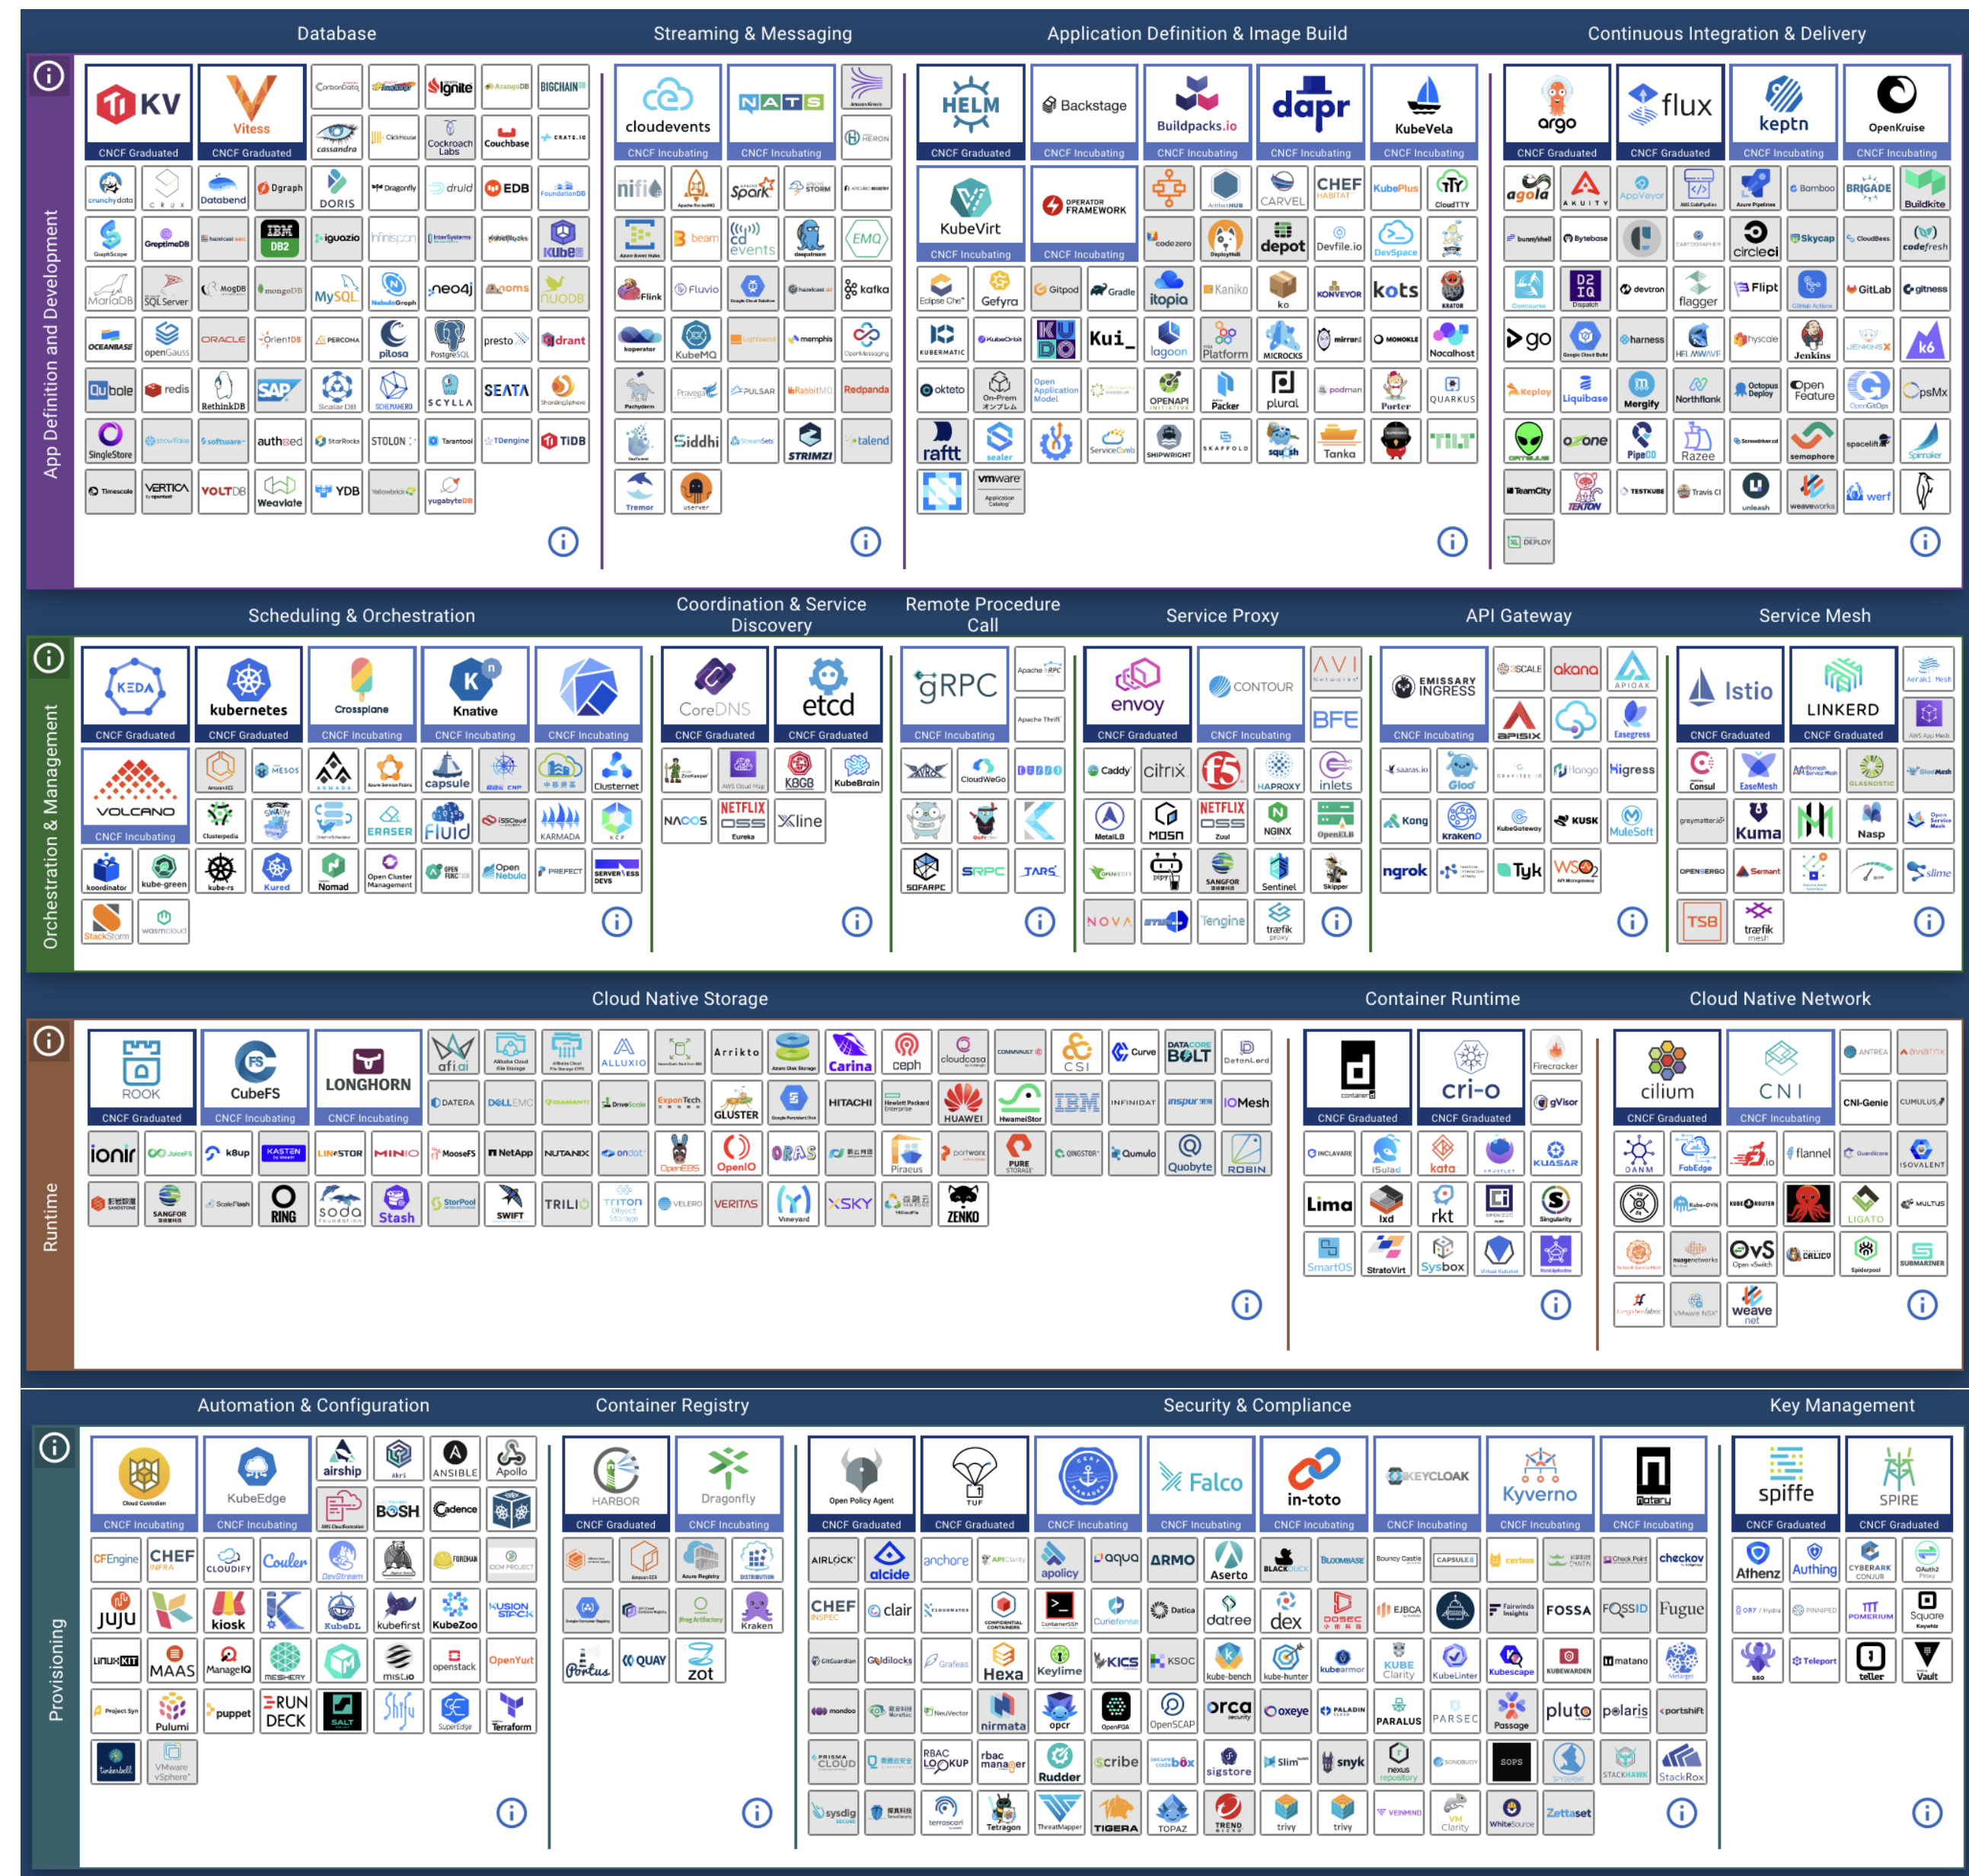
Surround Core Apps and DBs with New Microservices

2

Refactor Core Apps to New Microservices

3

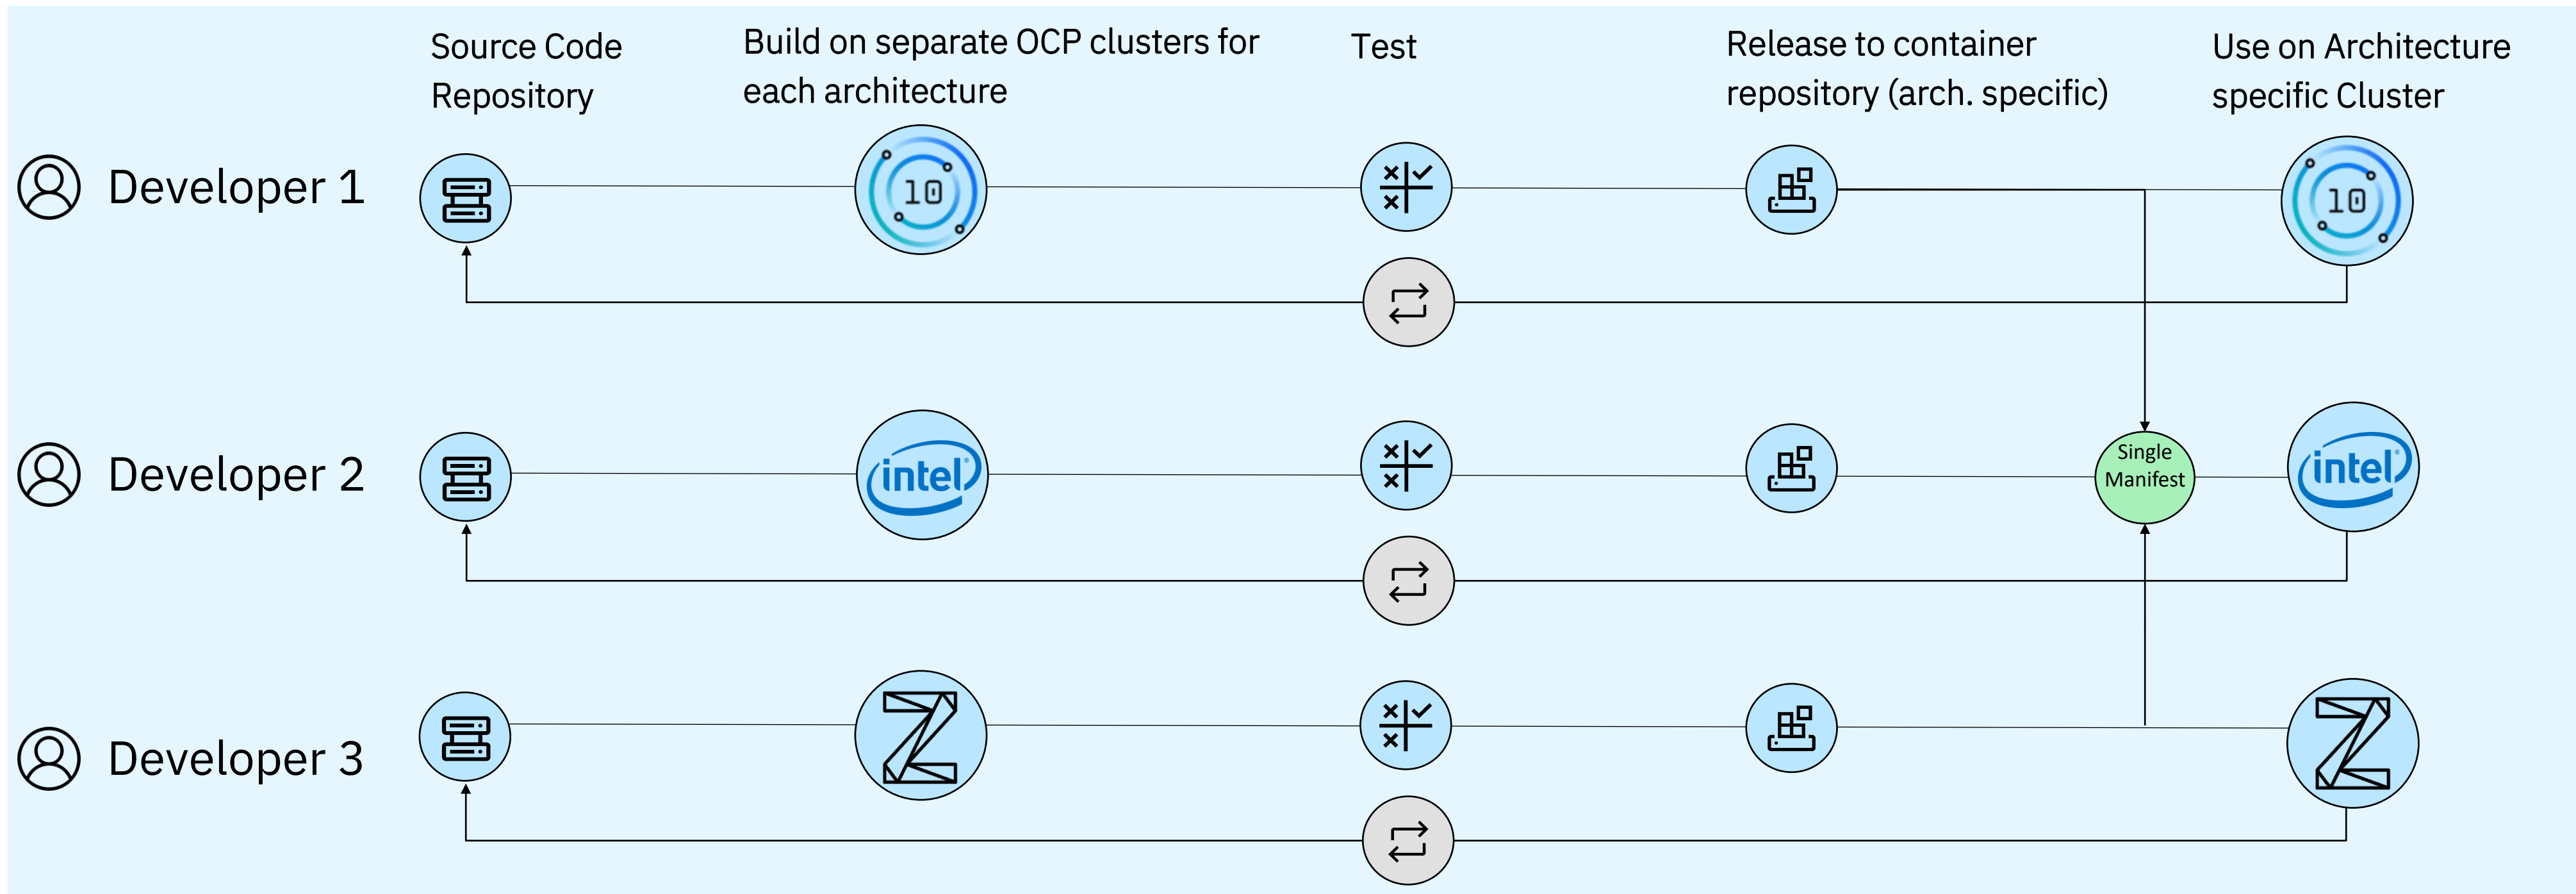
Refactor ISV or Custom Core Apps to Microservices



Multi-Arch Simplifies DevOps

Outcomes

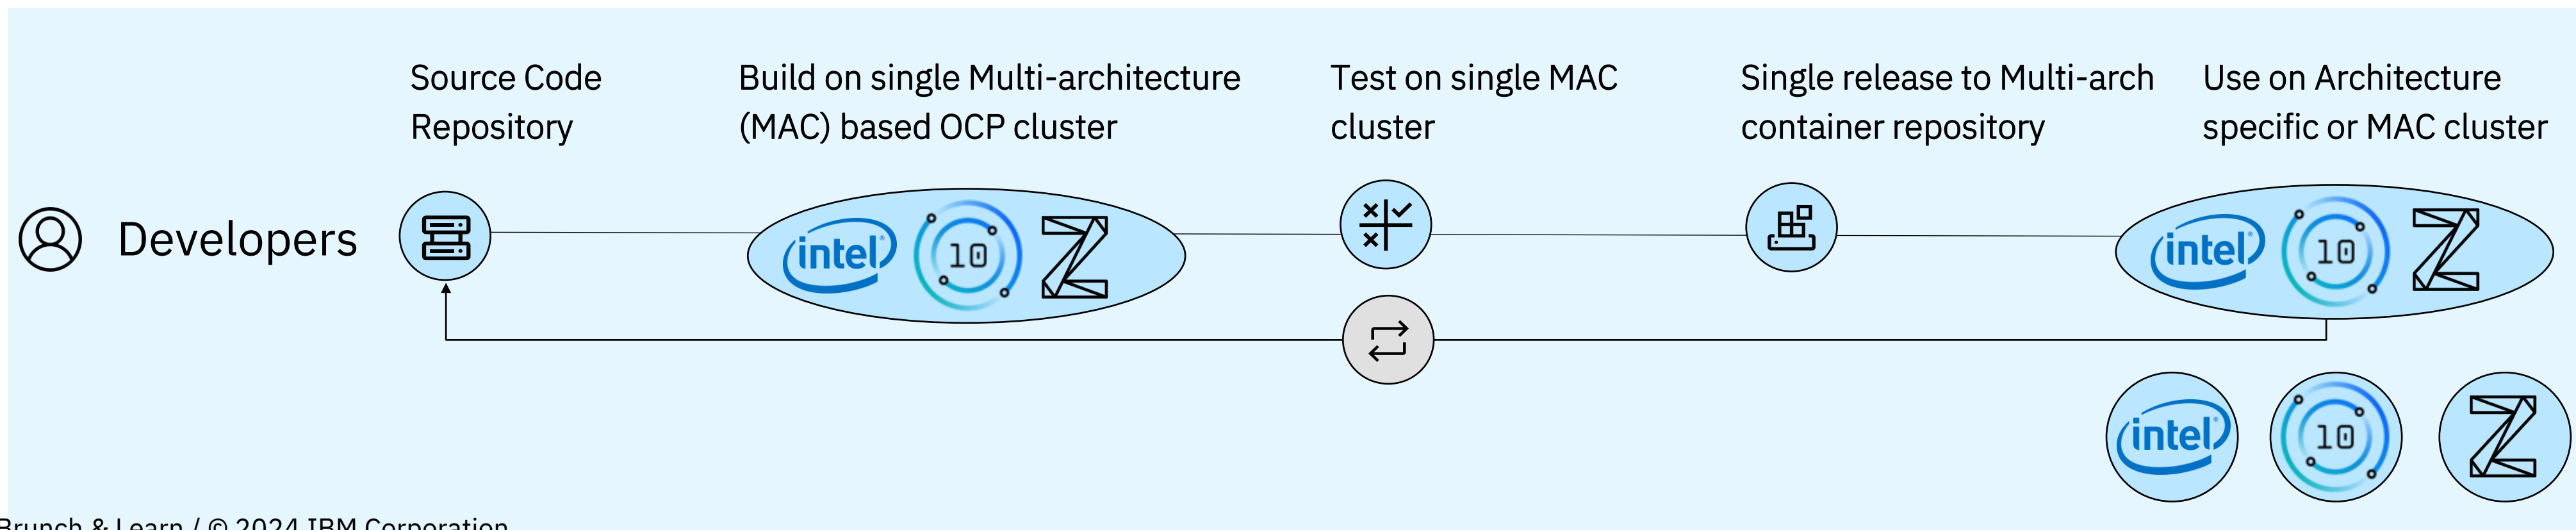
Before MAC



Before MAC

- Not as easy to be deterministic when there are 3 paths

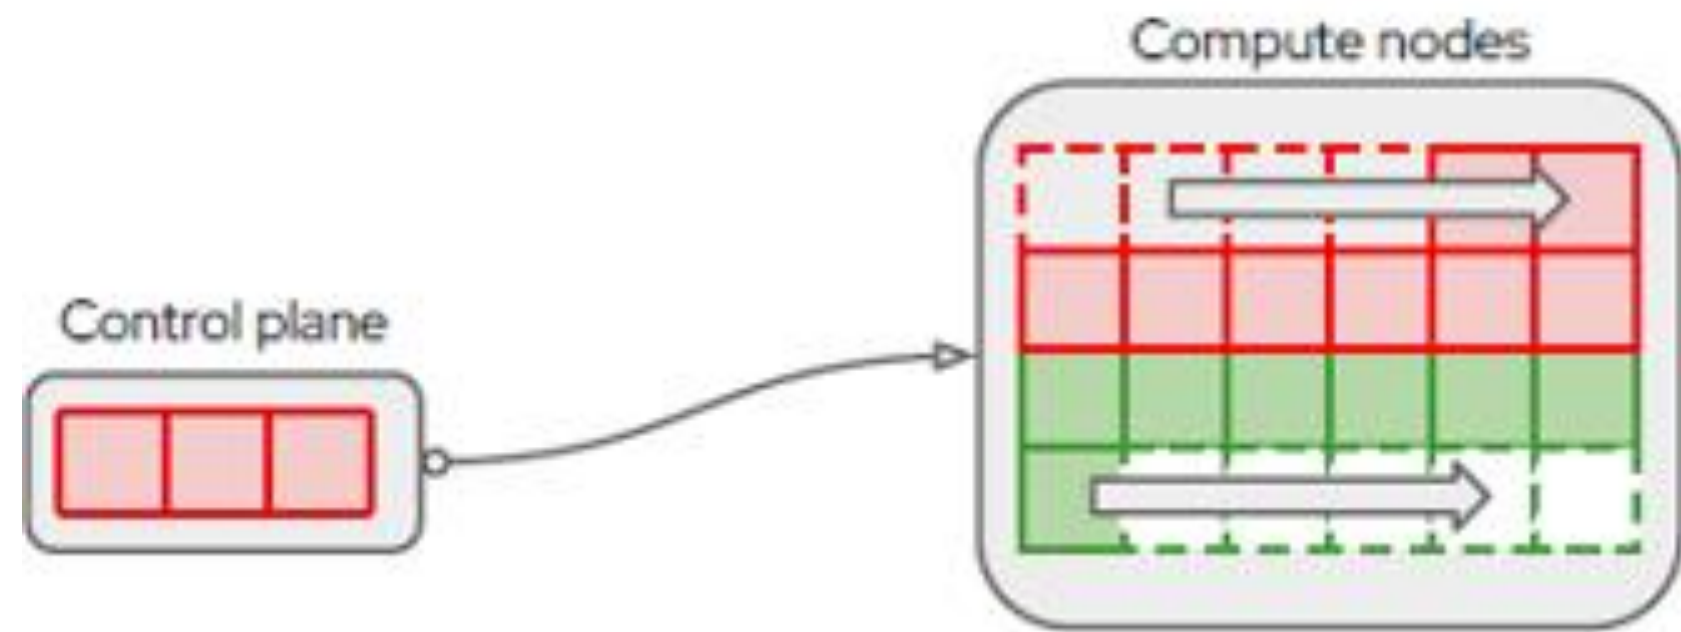
With MAC



With MAC

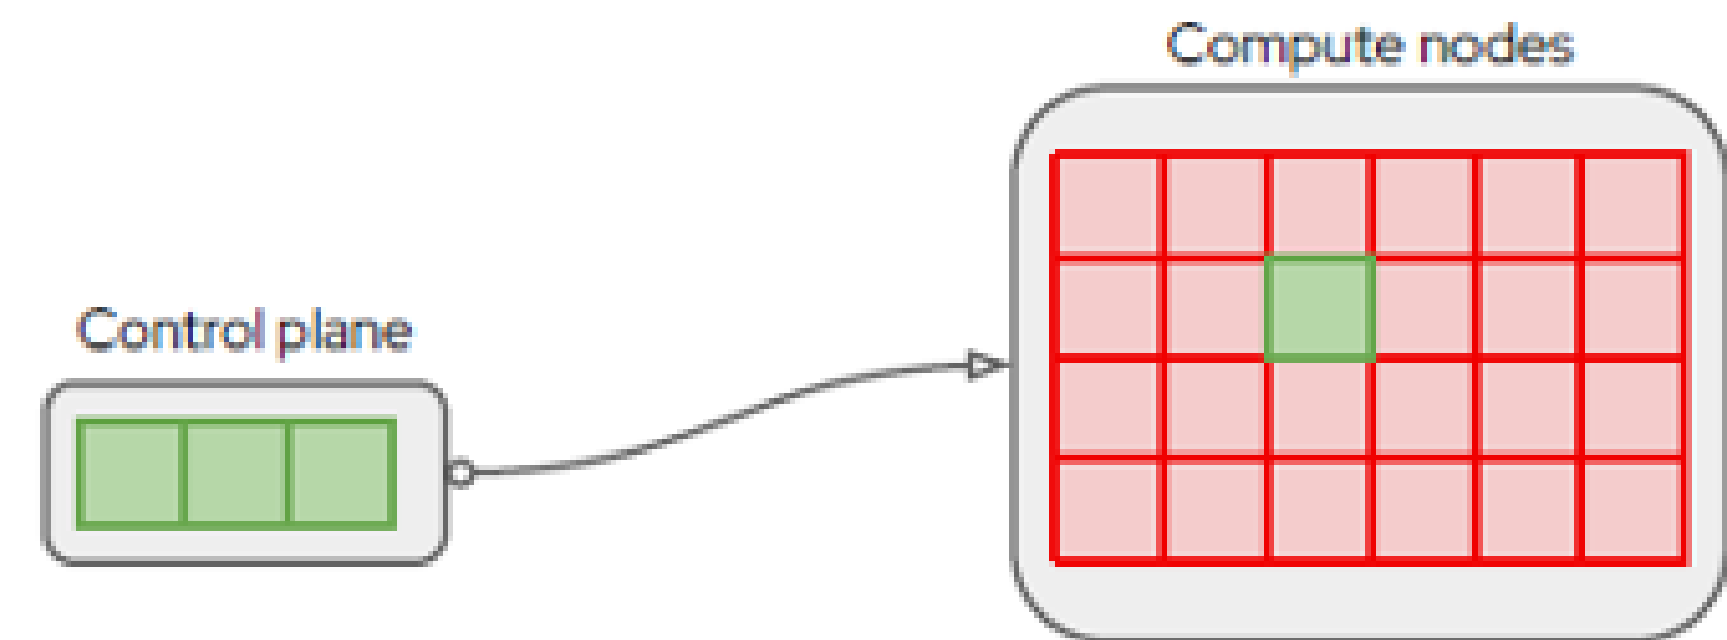
- Improved productivity (for devs, co-creation)
- Faster iterations which leads to faster releases and time to market
- Supports deterministic properties

Migration between Architectures



- Easy way to try out specific apps/services on a different architecture while maintaining your environment
- Gradual roll out of change to another architecture not requiring “big bang” change everything at once approach
- Interesting benchmark opportunity

Cost Optimisation



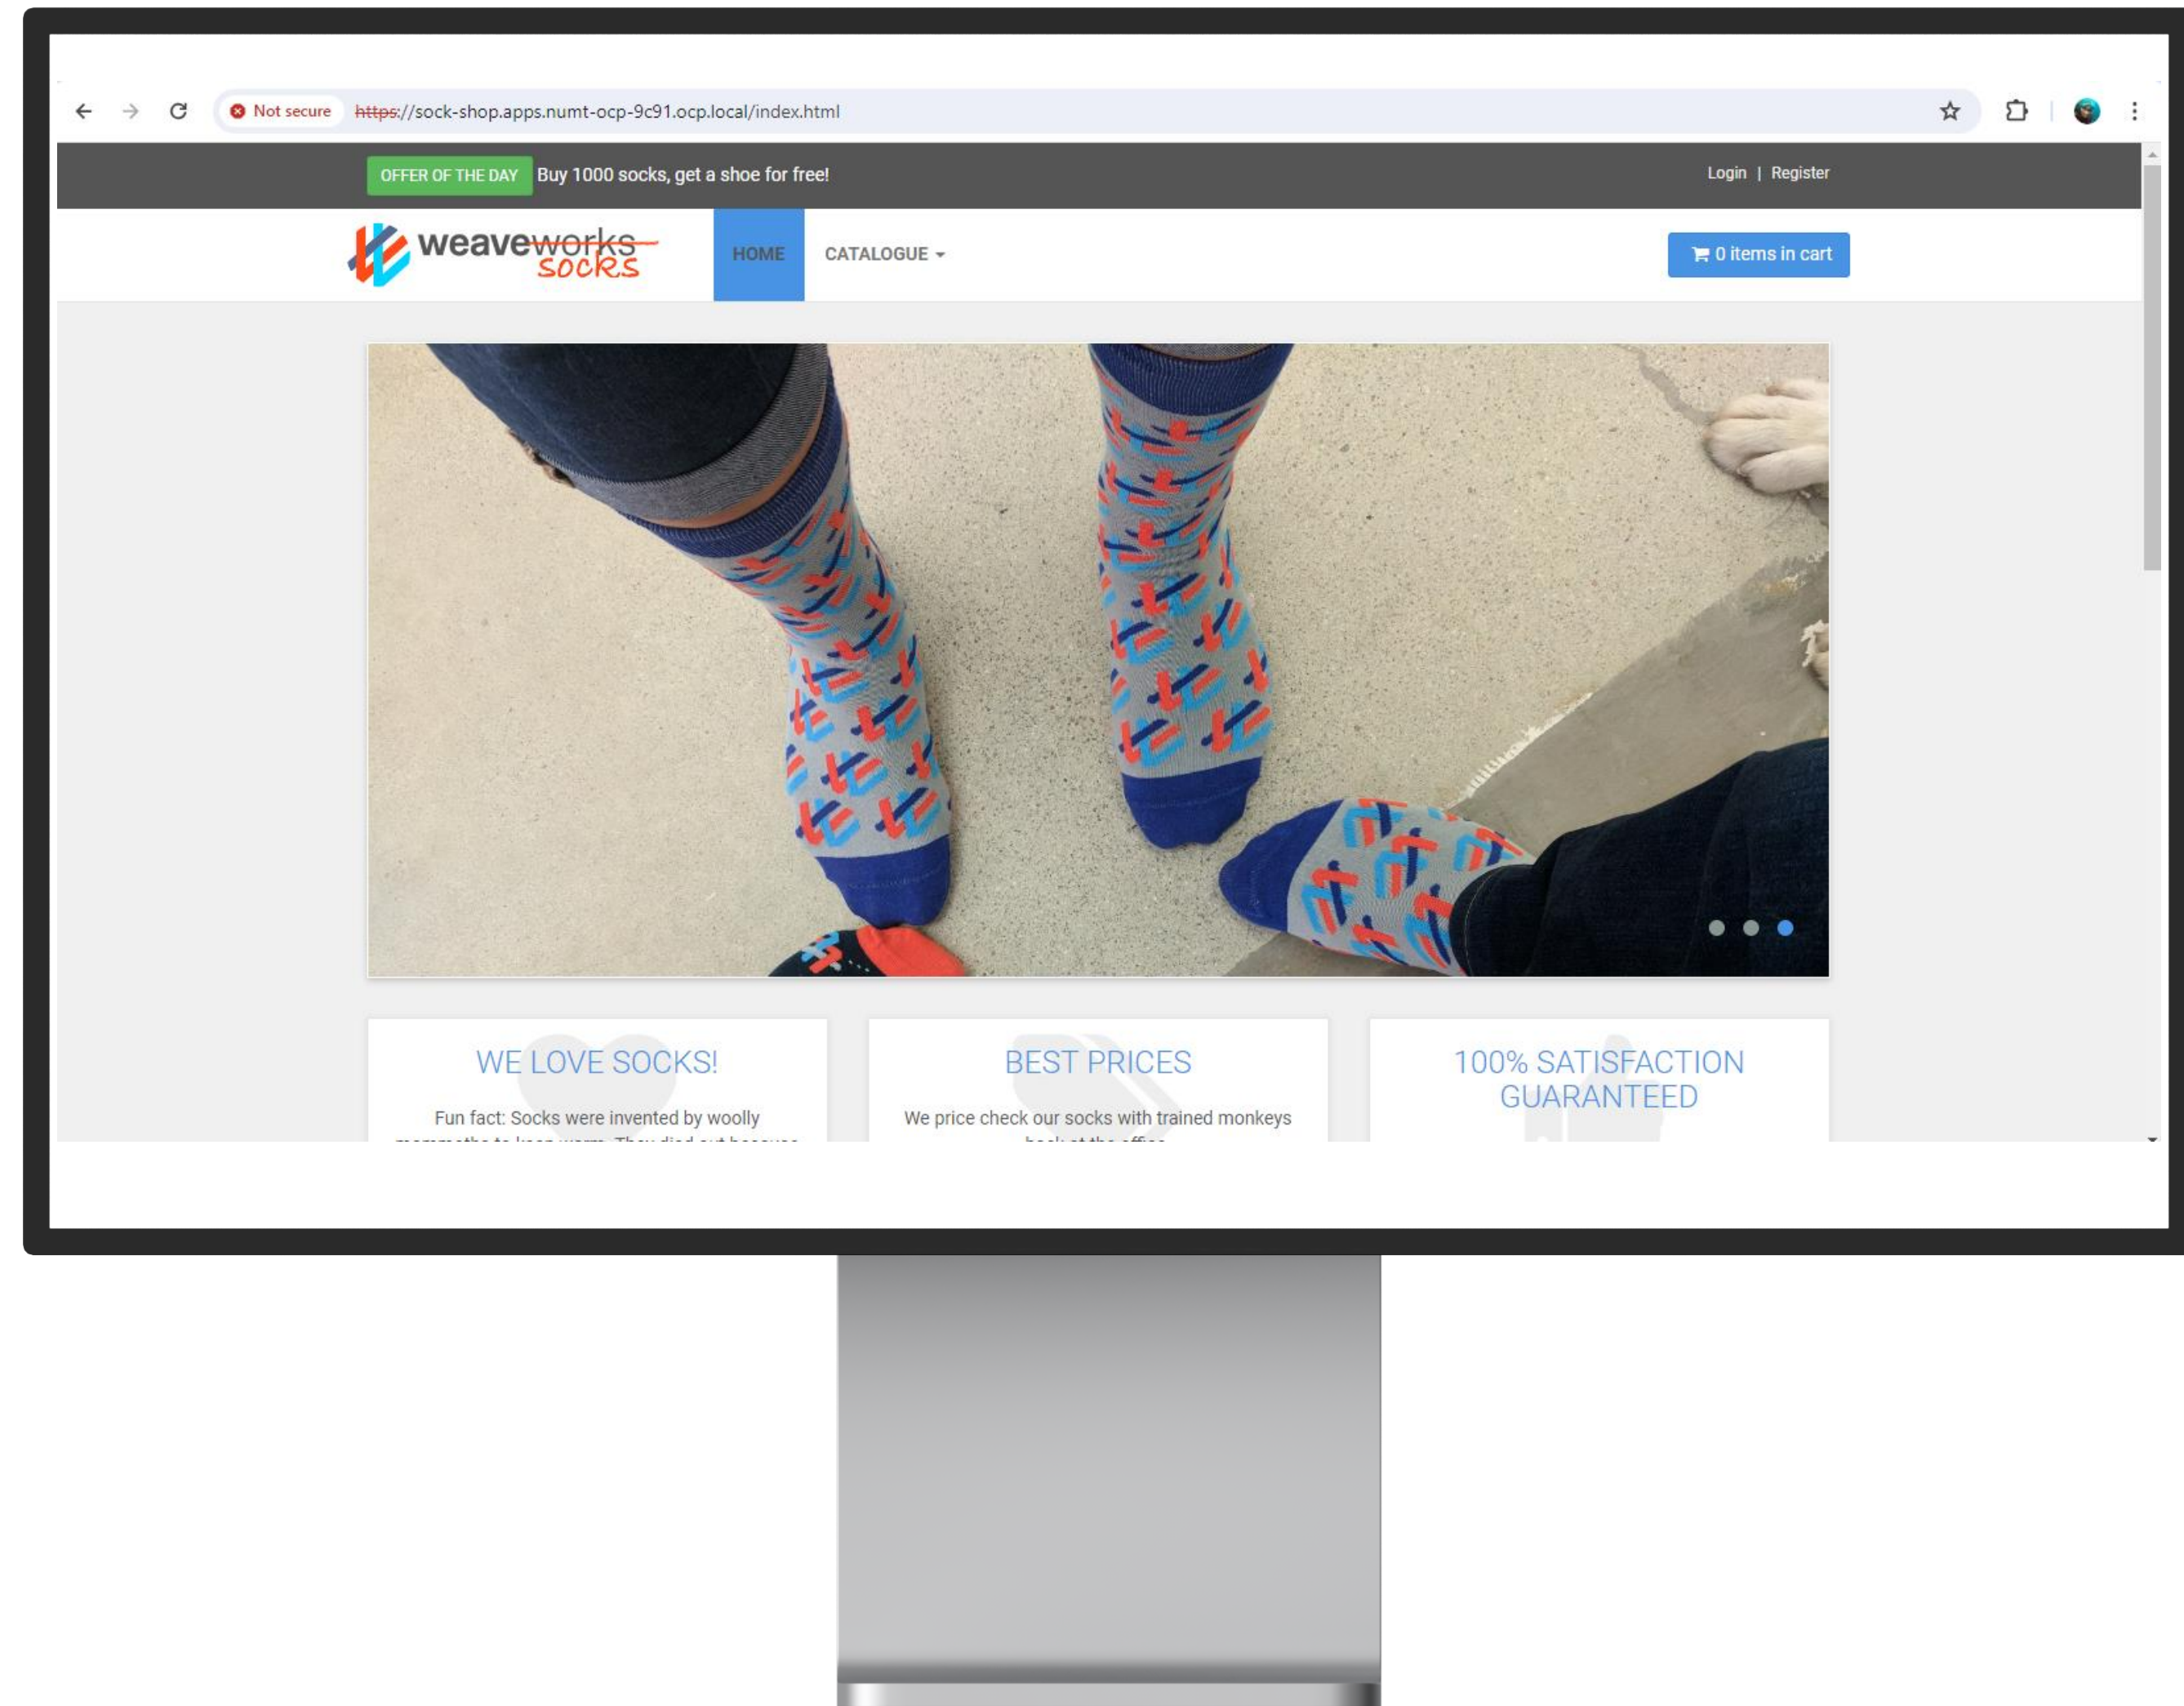
- Use a different architecture control plane that is made up of “cheaper” systems
- Interesting to cloud providers and also IBM zSystem and IBM Power

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 [Prepare for Installation](#)
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

Sock-Shop Website

- sock-shop front page screenshot
- This is what we are going to install



Available from GitHub

- I've forked the project to my repository in GitHub repository
- <https://github.com/paulchapmanibm/sock-shop-demo/tree/main>

The screenshot shows the GitHub interface for the repository `paulchapmanibm / sock-shop-demo`. The repository is public and was forked from `ocp-power-demos/sock-shop-demo`. The main branch is selected, and it is up to date with the upstream repository. The repository contains several files and folders:

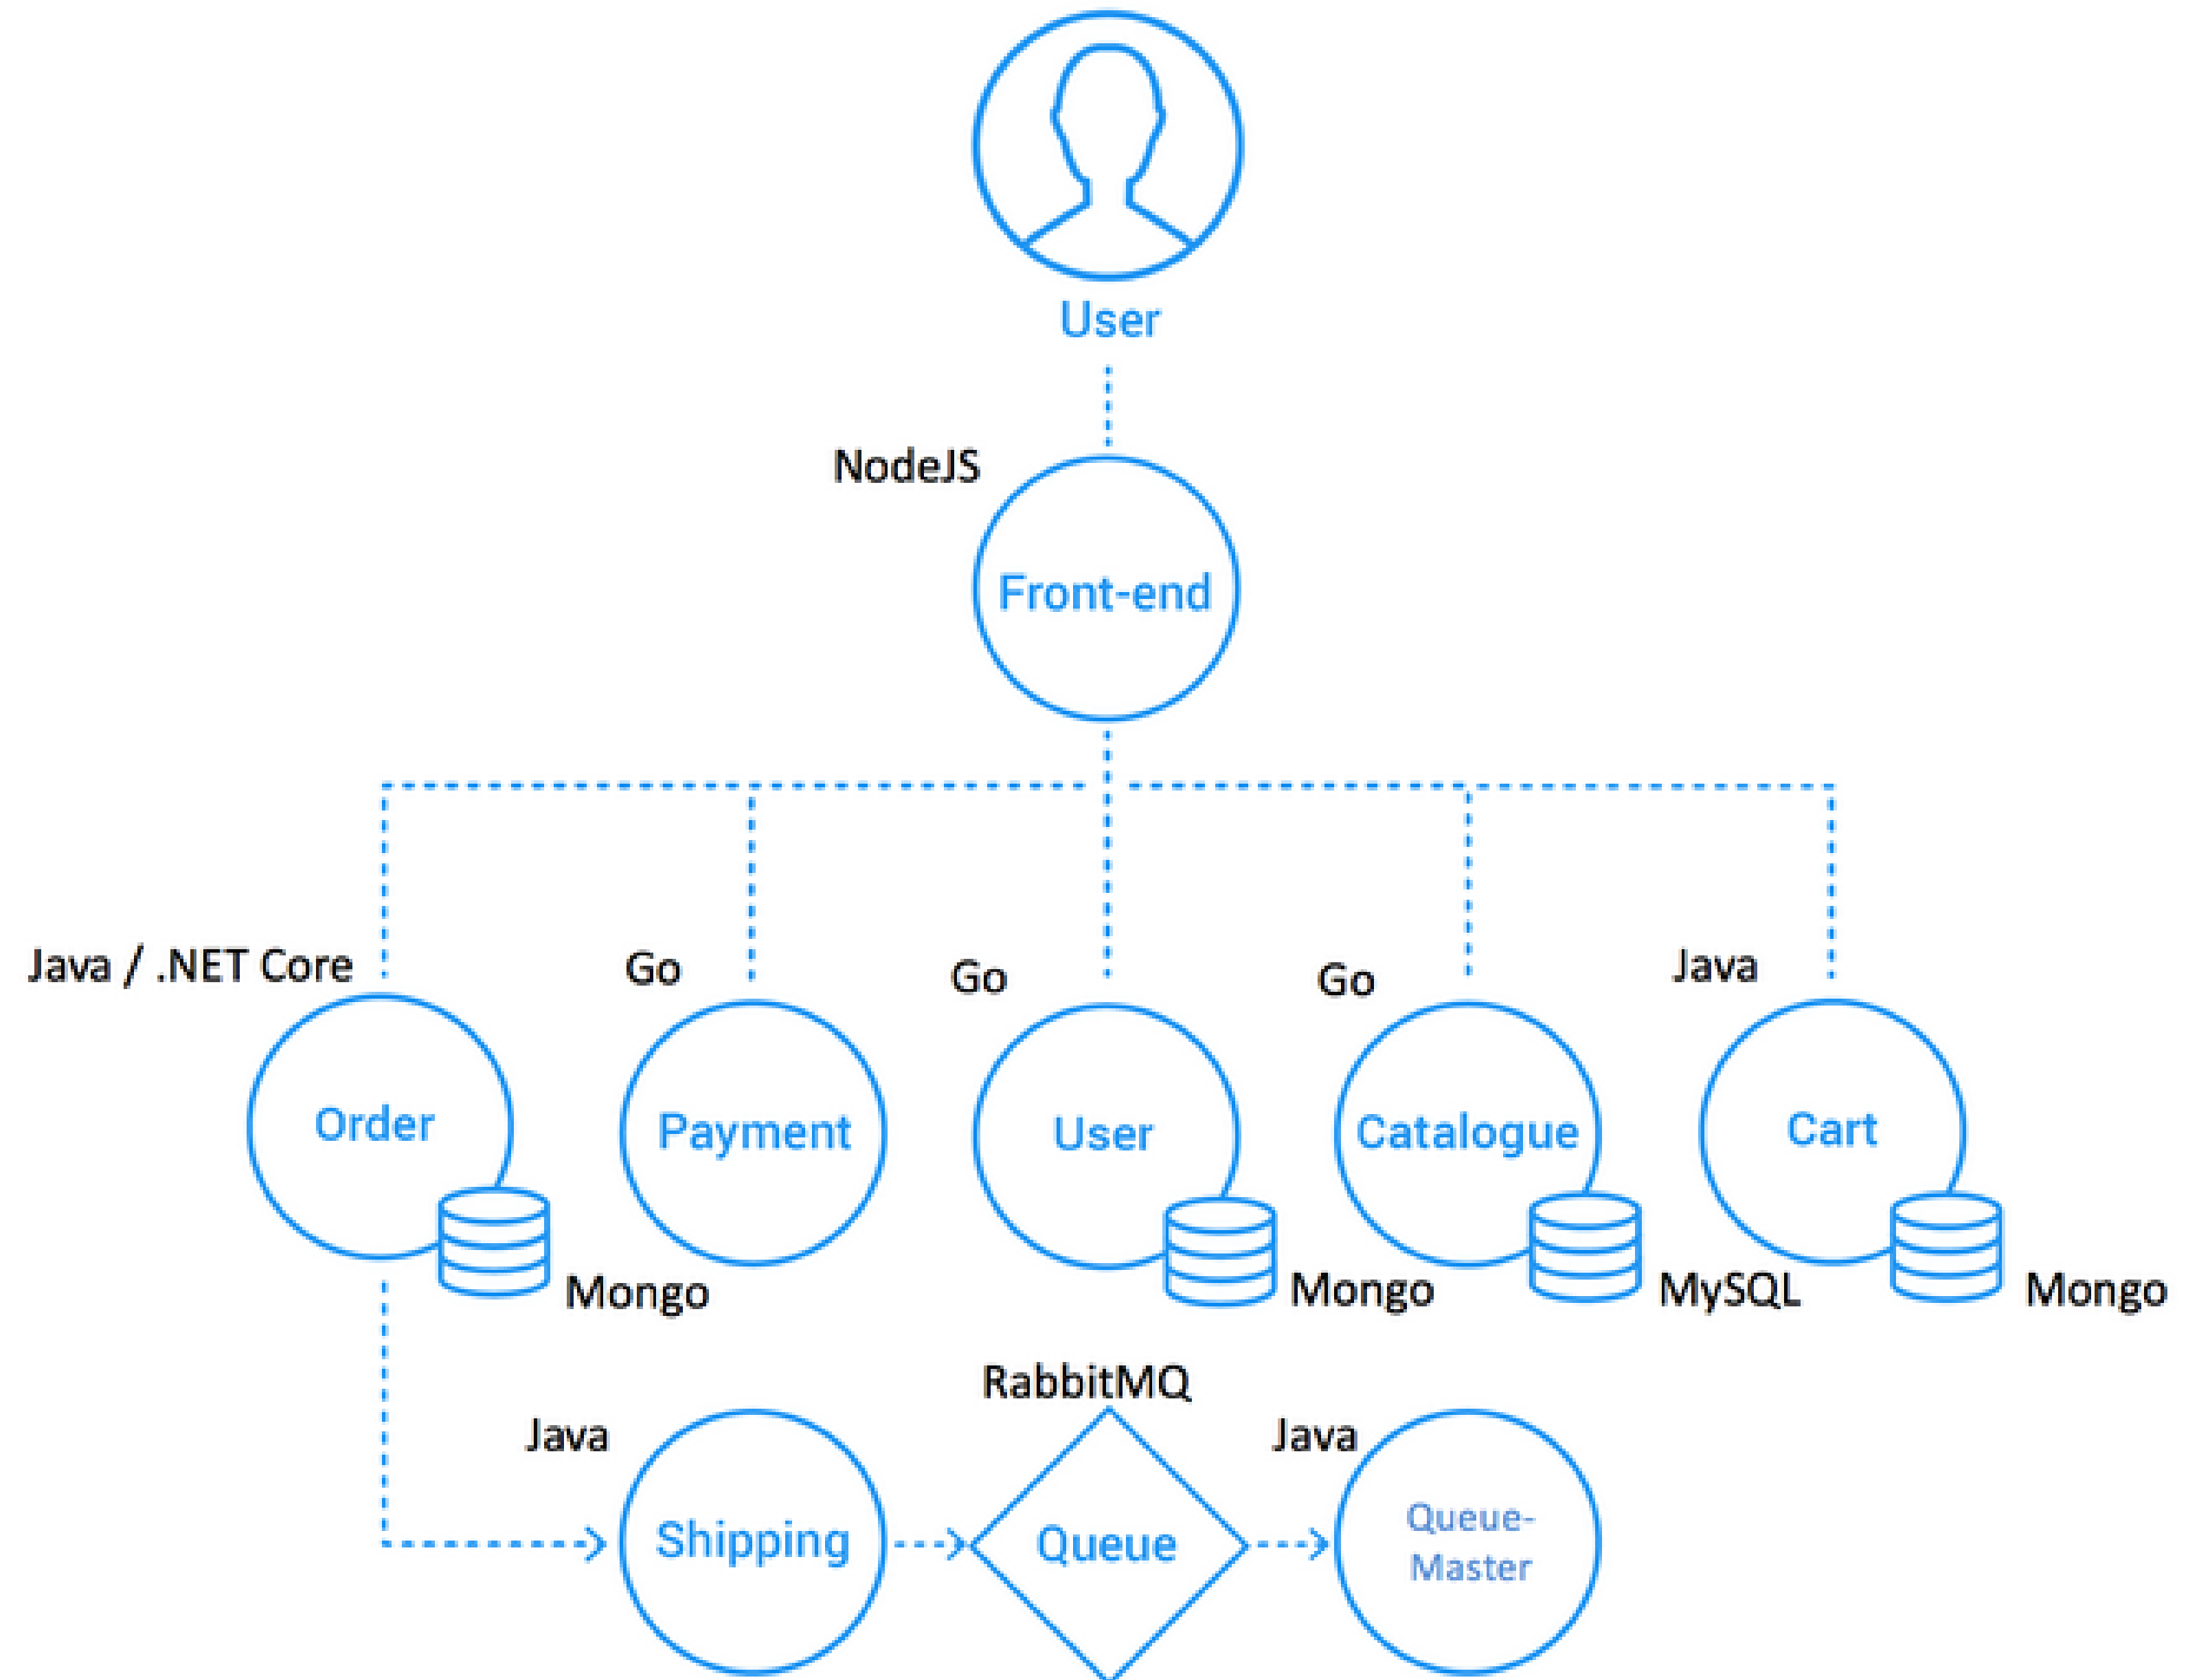
File/Folder	Commit Message	Commit Date
automation	Fix Power overlay	5 months ago
manifests	fix: missing storageclass in kustomization	5 days ago
.gitignore	migrate from prb112 to ocp-power-demos	9 months ago
LICENSE	Initial commit	9 months ago
Makefile	Updated Makefile for Multiarch image Support	8 months ago
README.md	Update README.md	6 months ago
socks-orders.png	migrate from prb112 to ocp-power-demos	9 months ago

The repository also includes a README, a License, and a list of commits. The commit history shows a recent commit by `prb112` with the message "fix: missing storageclass in kustomization" (note the typo) 5 days ago. The repository has 1 branch, 0 tags, and 28 commits. The repository statistics show 0 stars, 0 watching, and 1 fork.

Application Architecture

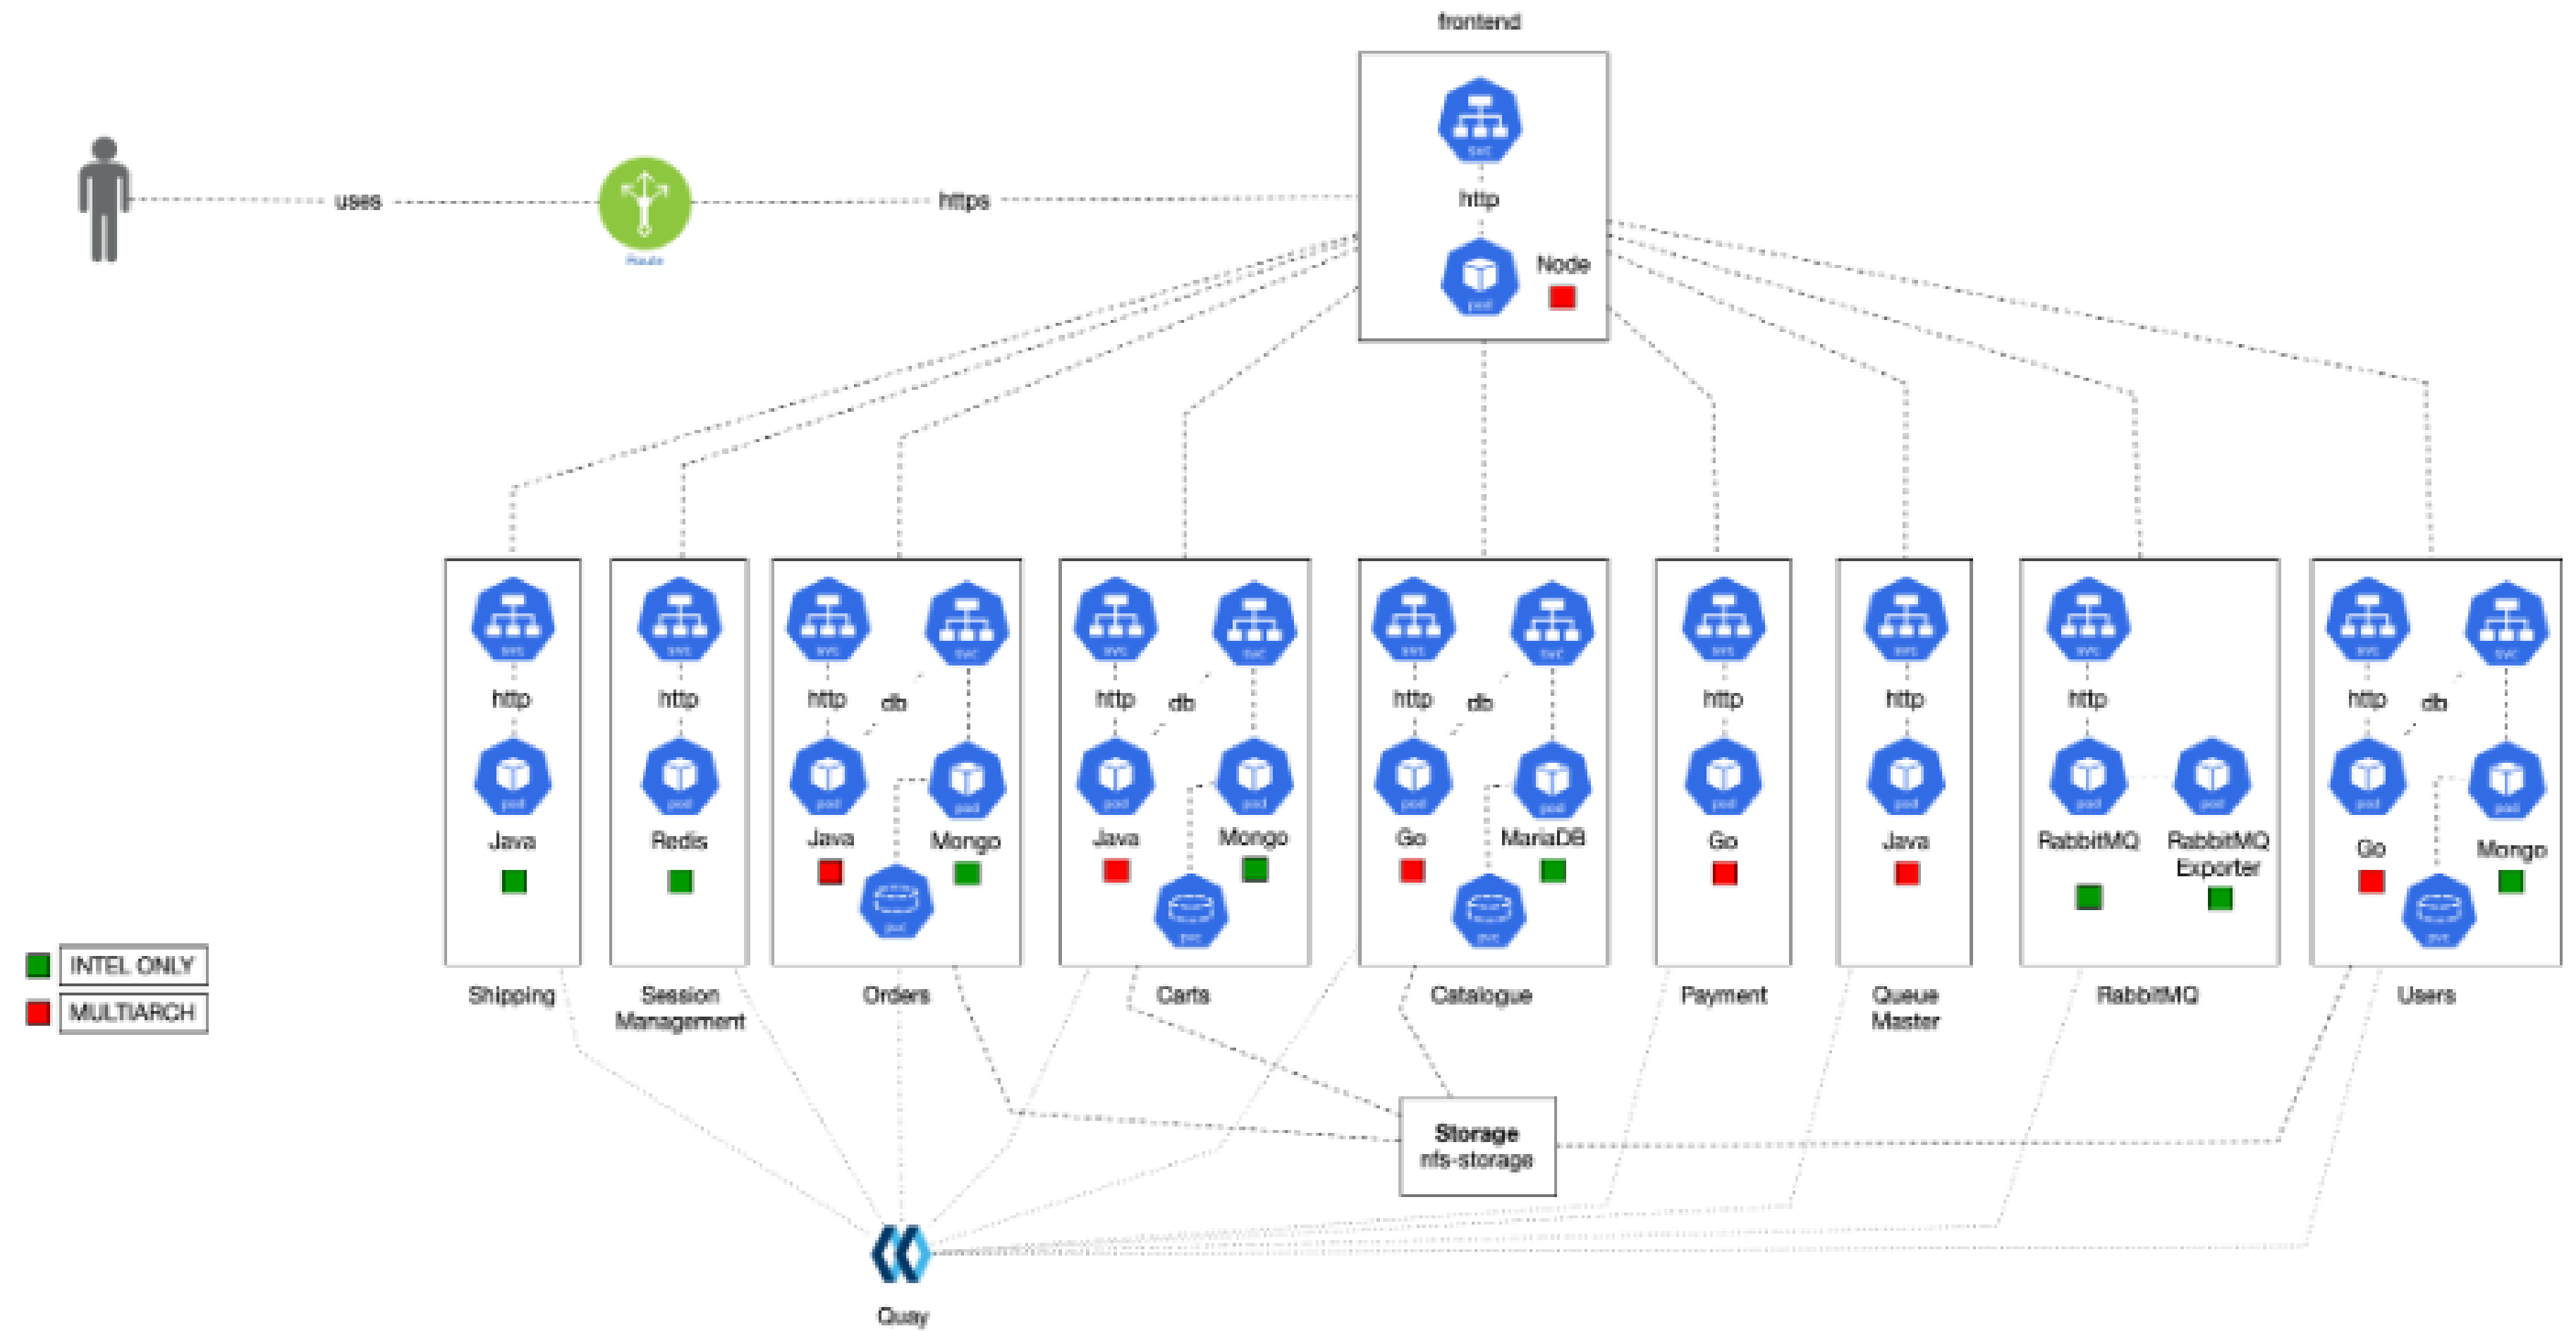
Applications

- front-end
- orders
- payment
- User-db is supported by power ppc64le overlay, otherwise is Intel only.
- catalogue catalogue-db is supported by Power ppc64le overlay, otherwise is Intel only.
- cart
- shipping
- queue-master



Microservice Connections

The microservices interaction diagram is:



- Users connect and log in via HTTPS

- NFS Storage
- Images are stored in Quay

Images from Quay

- Images are located in PowerCloud Quay repository
- <https://quay.io/organization/powercloud>
- Filter by sock-shop
- Note images for
 - ppc64le
 - arm64
 - s390x
 - amd64

The image shows a screenshot of the Quay.io web interface. At the top, the navigation bar includes 'RED HAT Quay.io', 'EXPLORE', 'REPOSITORIES', and 'TUTORIAL'. The user is logged in as 'paulchap...'. The main content area displays a list of repositories under the 'powercloud' organization, filtered by 'sock-shop'. The table below shows the details of these repositories:

REPOSITORY NAME	LAST MODIFIED	QUOTA CONSUMED	ACTIVITY ↓	STAR
powercloud / sock-shop-user	10/16/2023	368.28 MiB		☆
powercloud / sock-shop-orders	10/12/2023	808.81 MiB		☆
powercloud / sock-shop-shipping	10/12/2023	769.91 MiB		☆
powercloud / sock-shop-payment	10/12/2023	299.70 MiB		☆
powercloud / sock-shop-front-end	10/12/2023	1.42 GiB		☆
powercloud / sock-shop-catalogue	10/12/2023	633.77 MiB		☆
powercloud / sock-shop-carts	10/12/2023	677.02 MiB		☆
powercloud / sock-shop-catalogue-dl				
powercloud / sock-shop-user-db				
powercloud / sock-shop-queue-mast				

The bottom portion of the screenshot shows the 'Repository Tags' page for 'powercloud / sock-shop-payment'. It lists five tags with their respective details:

TAG	LAST MODIFIED ↓	SECURITY SCAN	SIZE	EXPIRES	MANIFEST
<input type="checkbox"/> ppc64le	8 months ago	14 High - 47 fixable	67.4 MiB	Never	SHA256 18515810f692
<input type="checkbox"/> arm64	8 months ago	14 High - 47 fixable	61.0 MiB	Never	SHA256 987ec8120593
<input type="checkbox"/> s390x	8 months ago	14 High - 47 fixable	106.5 MiB	Never	SHA256 ce2955abae55
<input type="checkbox"/> amd64	8 months ago	14 High - 47 fixable	64.7 MiB	Never	SHA256 4f9ffa82871f
<input type="checkbox"/> latest	8 months ago	See Child Manifests	N/A	Never	SHA256 e7c8a8538c82

Install Prereq's

- Install git
- Install Kustomize
- Install OC – At the same level as the OpenShift cluster
- 4.14 for x86 Control Plane
- 4.15 for Power Control Plane

```
paulc@li-31610ee6-5f8c-4ebb-bfff-5d760c6f7bf2:~  
File Edit View Search Terminal Help  
paulc:~$ sudo dnf install git  
Last metadata expiration check: 0:07:24 ago on Sun 02 Jun 2024 09:42:31 BST.  
Package git-2.39.3-1.el8_8.x86_64 is already installed.  
resolving  
Dependencies resolved.  
Nothing to do.  
Complete!  
paulc:~$  
paulc:~$  
paulc:~$ curl -s "https://raw.githubusercontent.com/kubernetes-sigs/kustomize/master/hack/install_kustomize.sh" | bash  
v5.4.2  
kustomize installed to /home/paulc/kustomize  
paulc:~$ sudo mv ./kustomize /usr/bin/  
paulc:~$  
paulc:~$  
paulc:~$ which oc  
/usr/bin/oc  
paulc:~$ which kustomize  
/usr/bin/kustomize  
paulc:~$ █
```

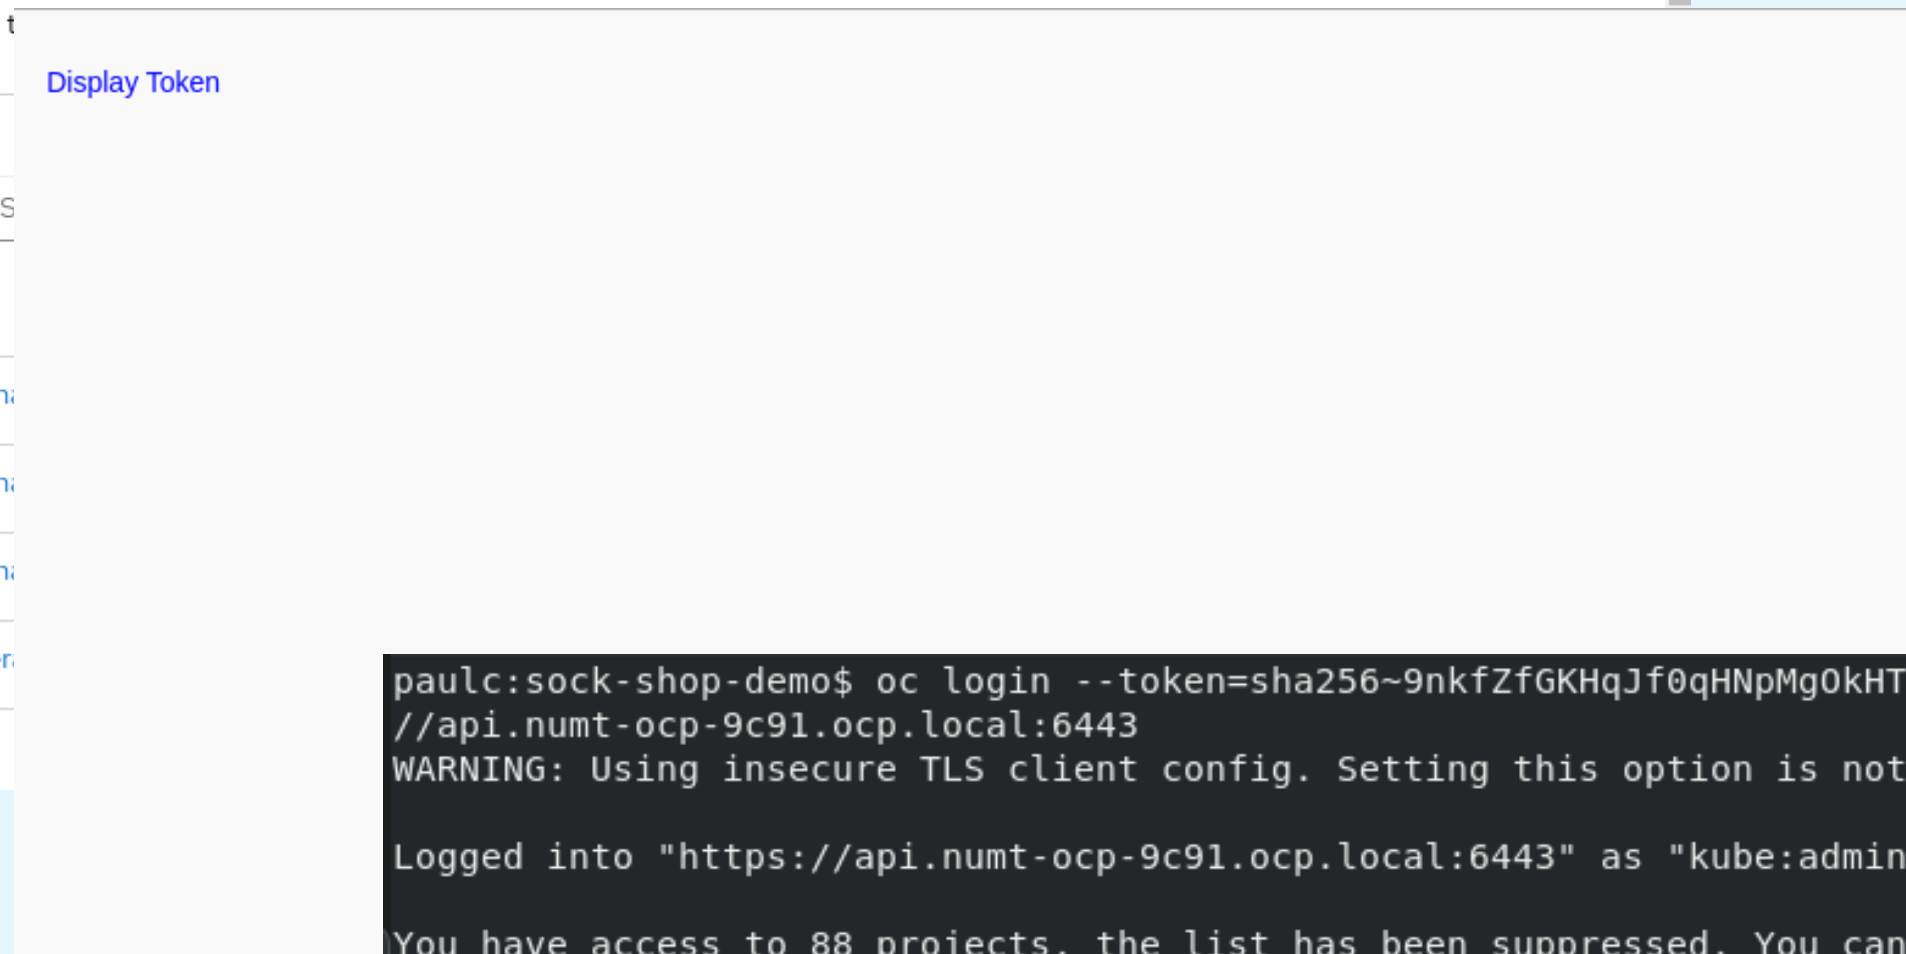
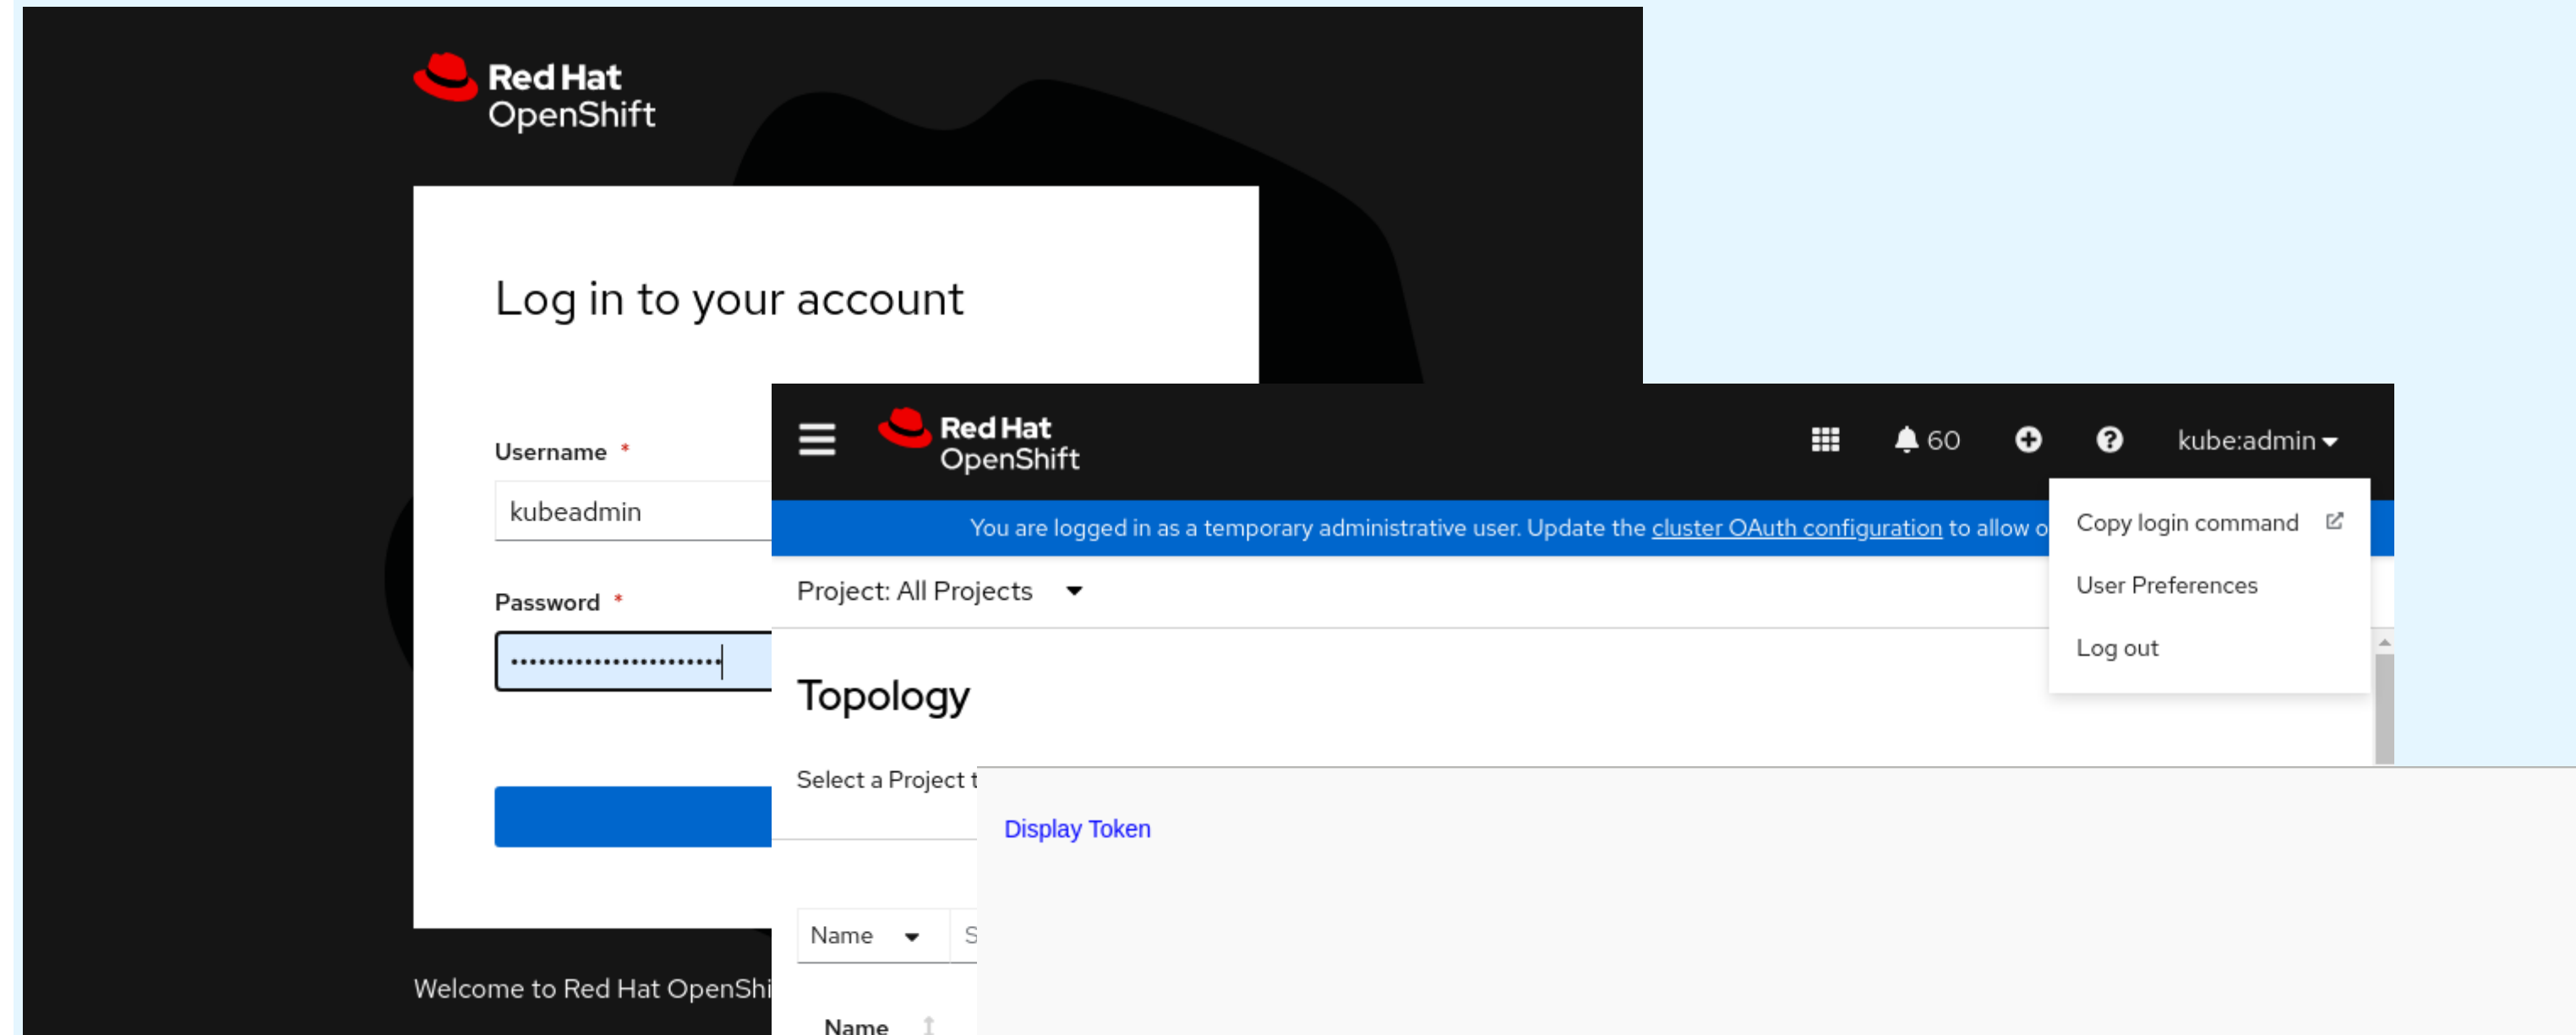

Clone & Configure Project

- Clone sock-shop project from GitHub
- Add application secrets

```
File Edit View Search Terminal Help
paulc:~$ pwd
/home/paulc
paulc:~$ git clone https://github.com/paulchapmanibm/sock-shop-demo.git
Cloning into 'sock-shop-demo'...
remote: Enumerating objects: 307, done.
remote: Counting objects: 100% (307/307), done.
remote: Compressing objects: 100% (165/165), done.
remote: Total 307 (delta 218), reused 197 (delta 138), pack-reused 0
Receiving objects: 100% (307/307), 152.75 KiB | 1.15 MiB/s, done.
Resolving deltas: 100% (218/218), done.
paulc:~$
paulc:~$ cp ~/env.secret ~/sock-shop-demo/manifests/base/env.secret
paulc:~$ █
```

CLI Login

- Log into OCP Console
- Select copy login command
- Select Display Token
- Copy & Paste OC login command into CLI



Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 [Install Multi-Arch sock-shop](#)
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

Deploy Solution

- Change to sock-shop-demo directory
- Deploy sock-shop application to OpenShift

```
File Edit View Search Terminal Help
paulc:~$ cd sock-shop-demo
paulc:sock-shop-demo$ kustomize build manifests/overlays/multi | oc apply -f -
project.project.openshift.io/sock-shop created
deployment.apps/carts created
service/carts created
persistentvolumeclaim/carts-db-temp-pvc created
deployment.apps/carts-db created
service/carts-db created
deployment.apps/catalogue created
service/catalogue created
deployment.apps/catalogue-db created
service/catalogue-db created
deployment.apps/front-end created
service/front-end created
deployment.apps/orders created
service/orders created
persistentvolumeclaim/orders-db-temp-pvc created
deployment.apps/orders-db created
service/orders-db created
deployment.apps/payment created
service/payment created
deployment.apps/queue-master created
service/queue-master created
deployment.apps/rabbitmq created
service/rabbitmq created
deployment.apps/session-db created
service/session-db created
deployment.apps/shipping created
service/shipping created
deployment.apps/user created
service/user created
persistentvolumeclaim/user-db-temp-pvc created
deployment.apps/user-db created
service/user-db created
route.route.openshift.io/sock-shop created
secret/monqodb-creds created
```

View Topology

- Select Topology from OCP Console
- Note that dark blue circles indicate that Pod is stable. Light blue indicates building, and red a problem

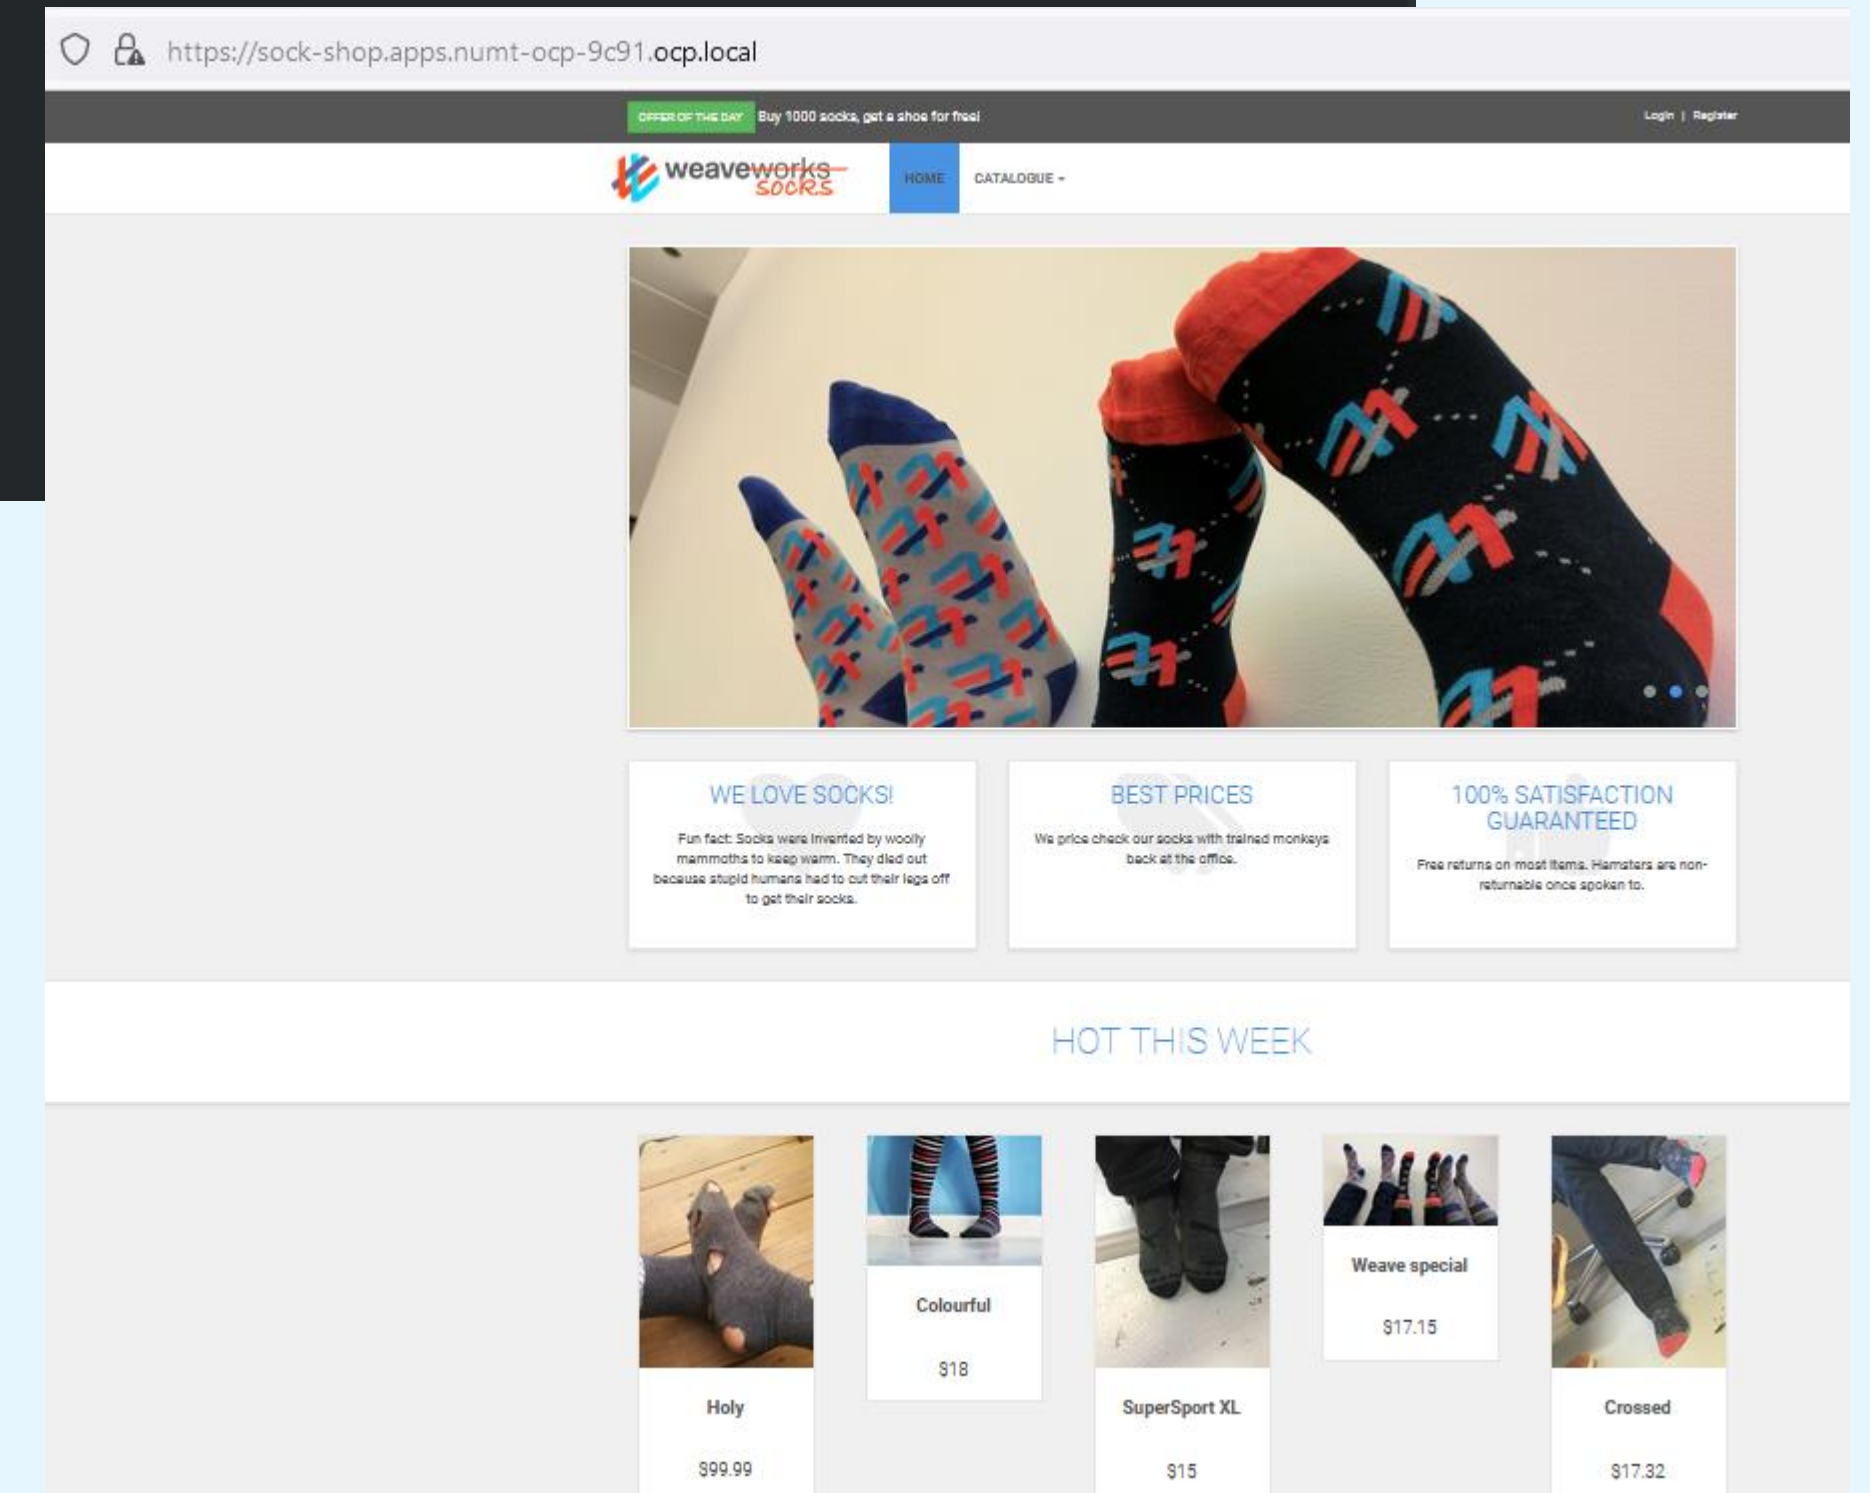
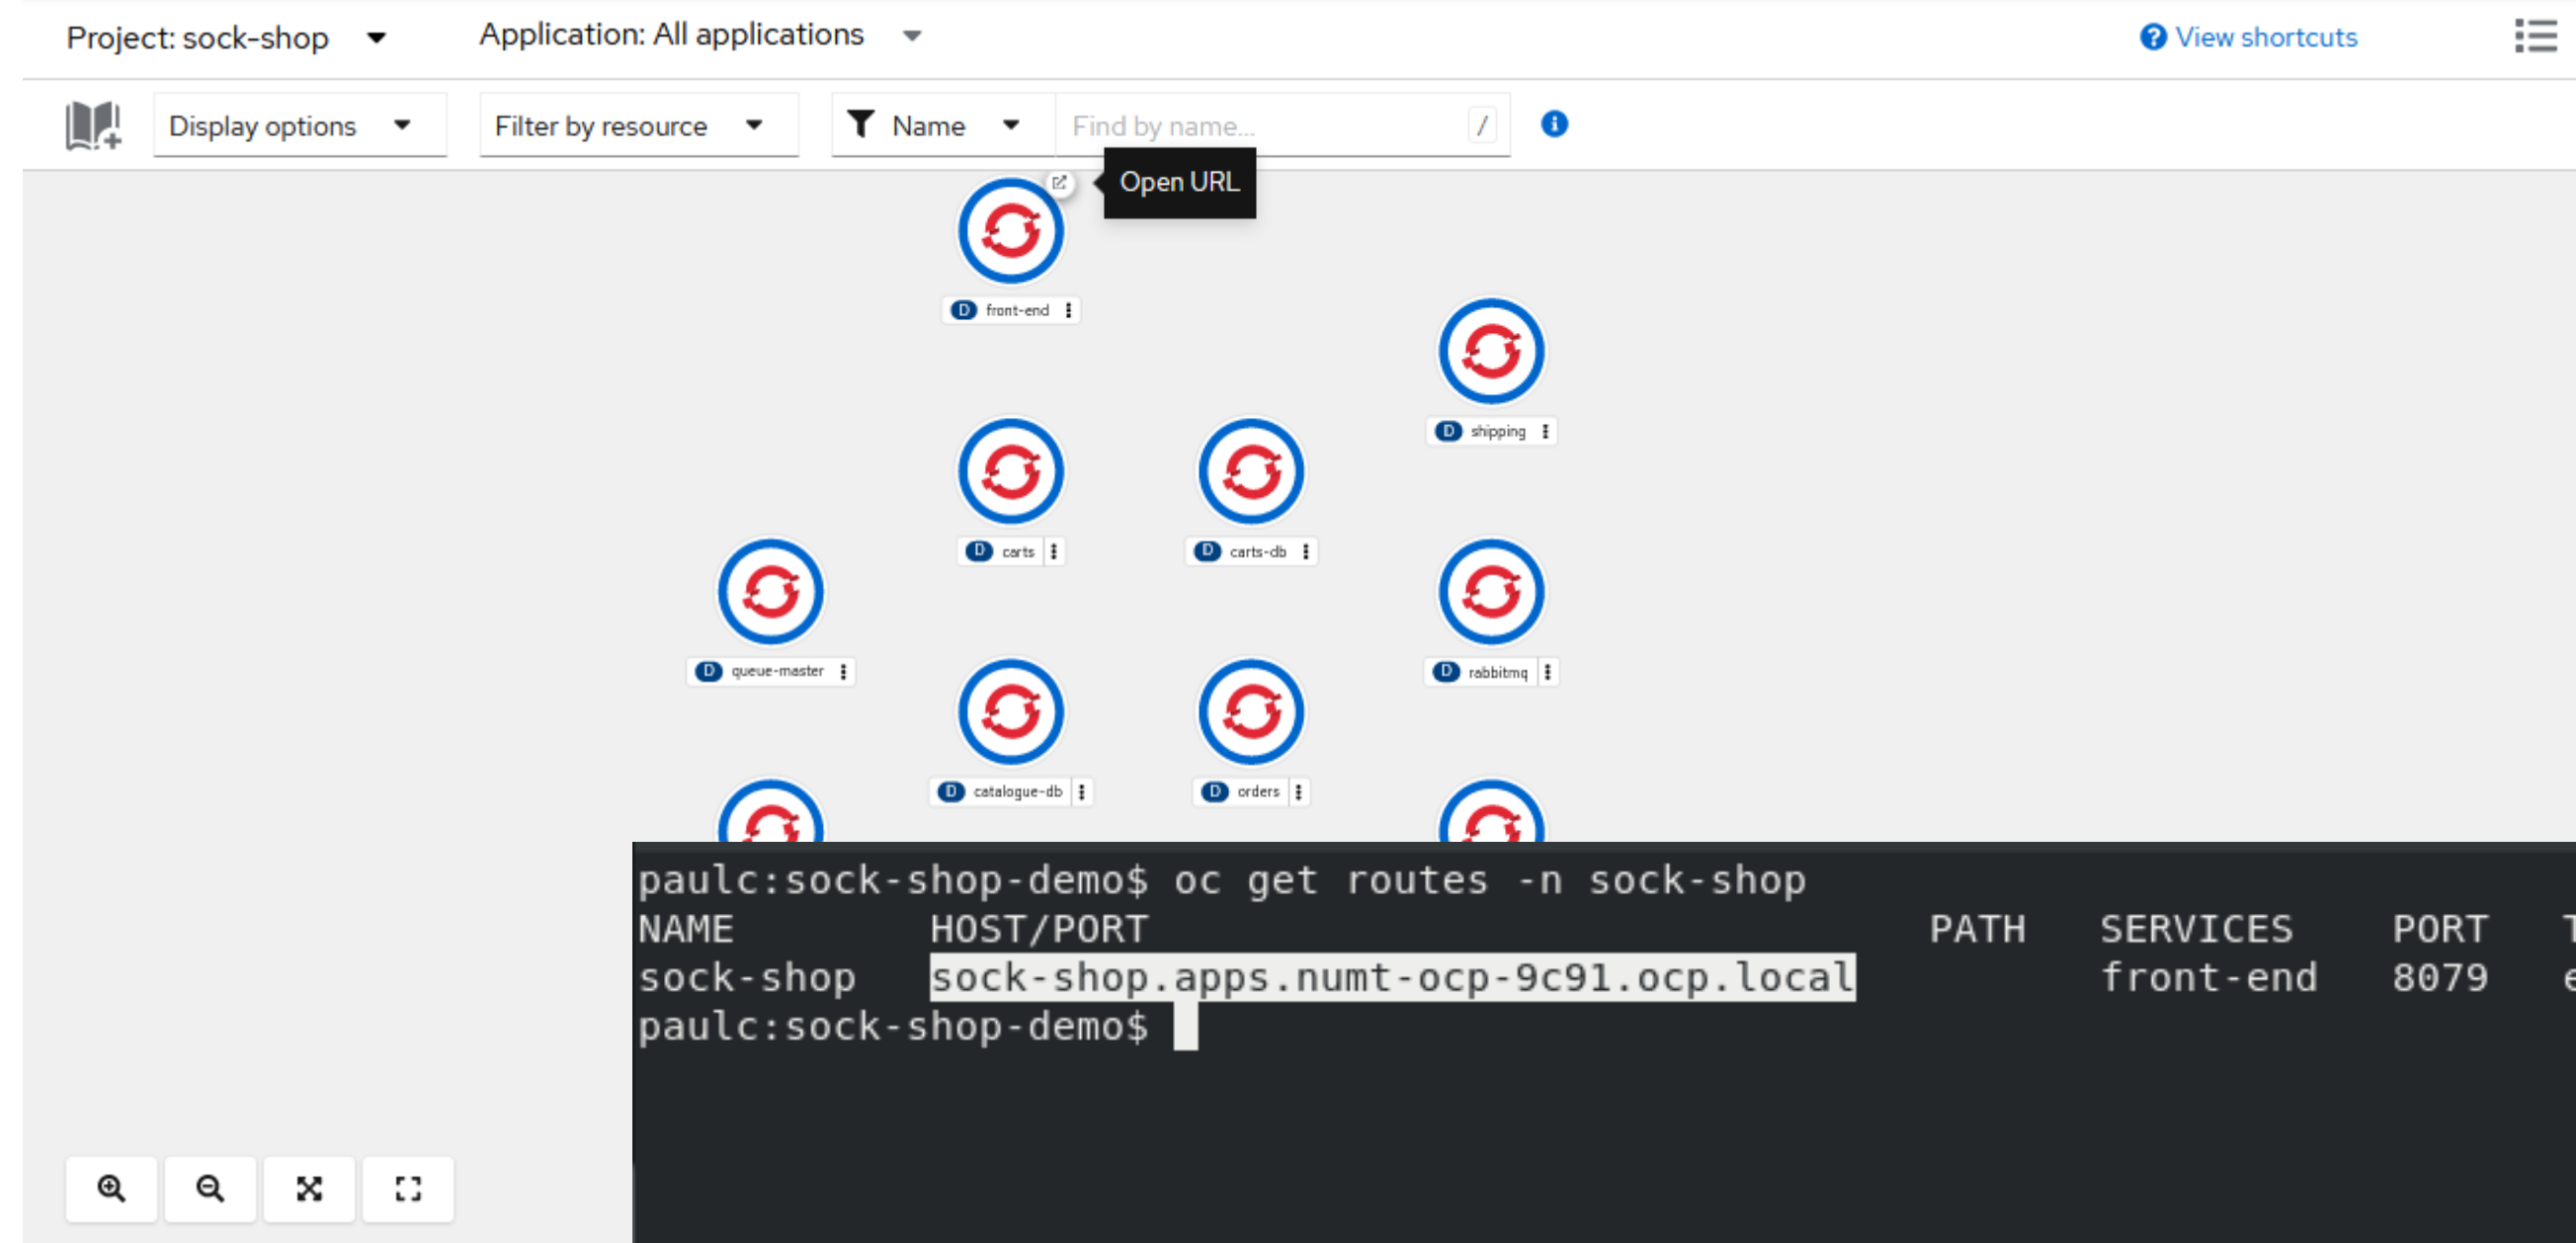
The screenshot shows the OpenShift OCP Console interface for the 'sock-shop' project. The top navigation bar indicates the user is logged in as a temporary administrative user. The main content area displays a topology view of the application's pods. The pods are arranged in a grid, and their status is indicated by the color of the pod icon: dark blue for stable, light blue for building, and red for a problem. The 'front-end' pod is stable (dark blue), while the 'carts', 'carts-db', 'shipping', 'rabbitmq', 'session-db', and 'orders-db' pods are in a red state, indicating a problem. The 'queue-master', 'catalogue-db', 'orders', 'user', 'user-db', and 'payment' pods are in a light blue state, indicating they are building. The 'orders-db' pod at the bottom is also in a red state. The left sidebar shows the navigation menu with 'Topology' selected. The top right corner has a 'View shortcuts' link. The bottom left corner has search and zoom controls.

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 [Navigate the sock-shop](#)
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

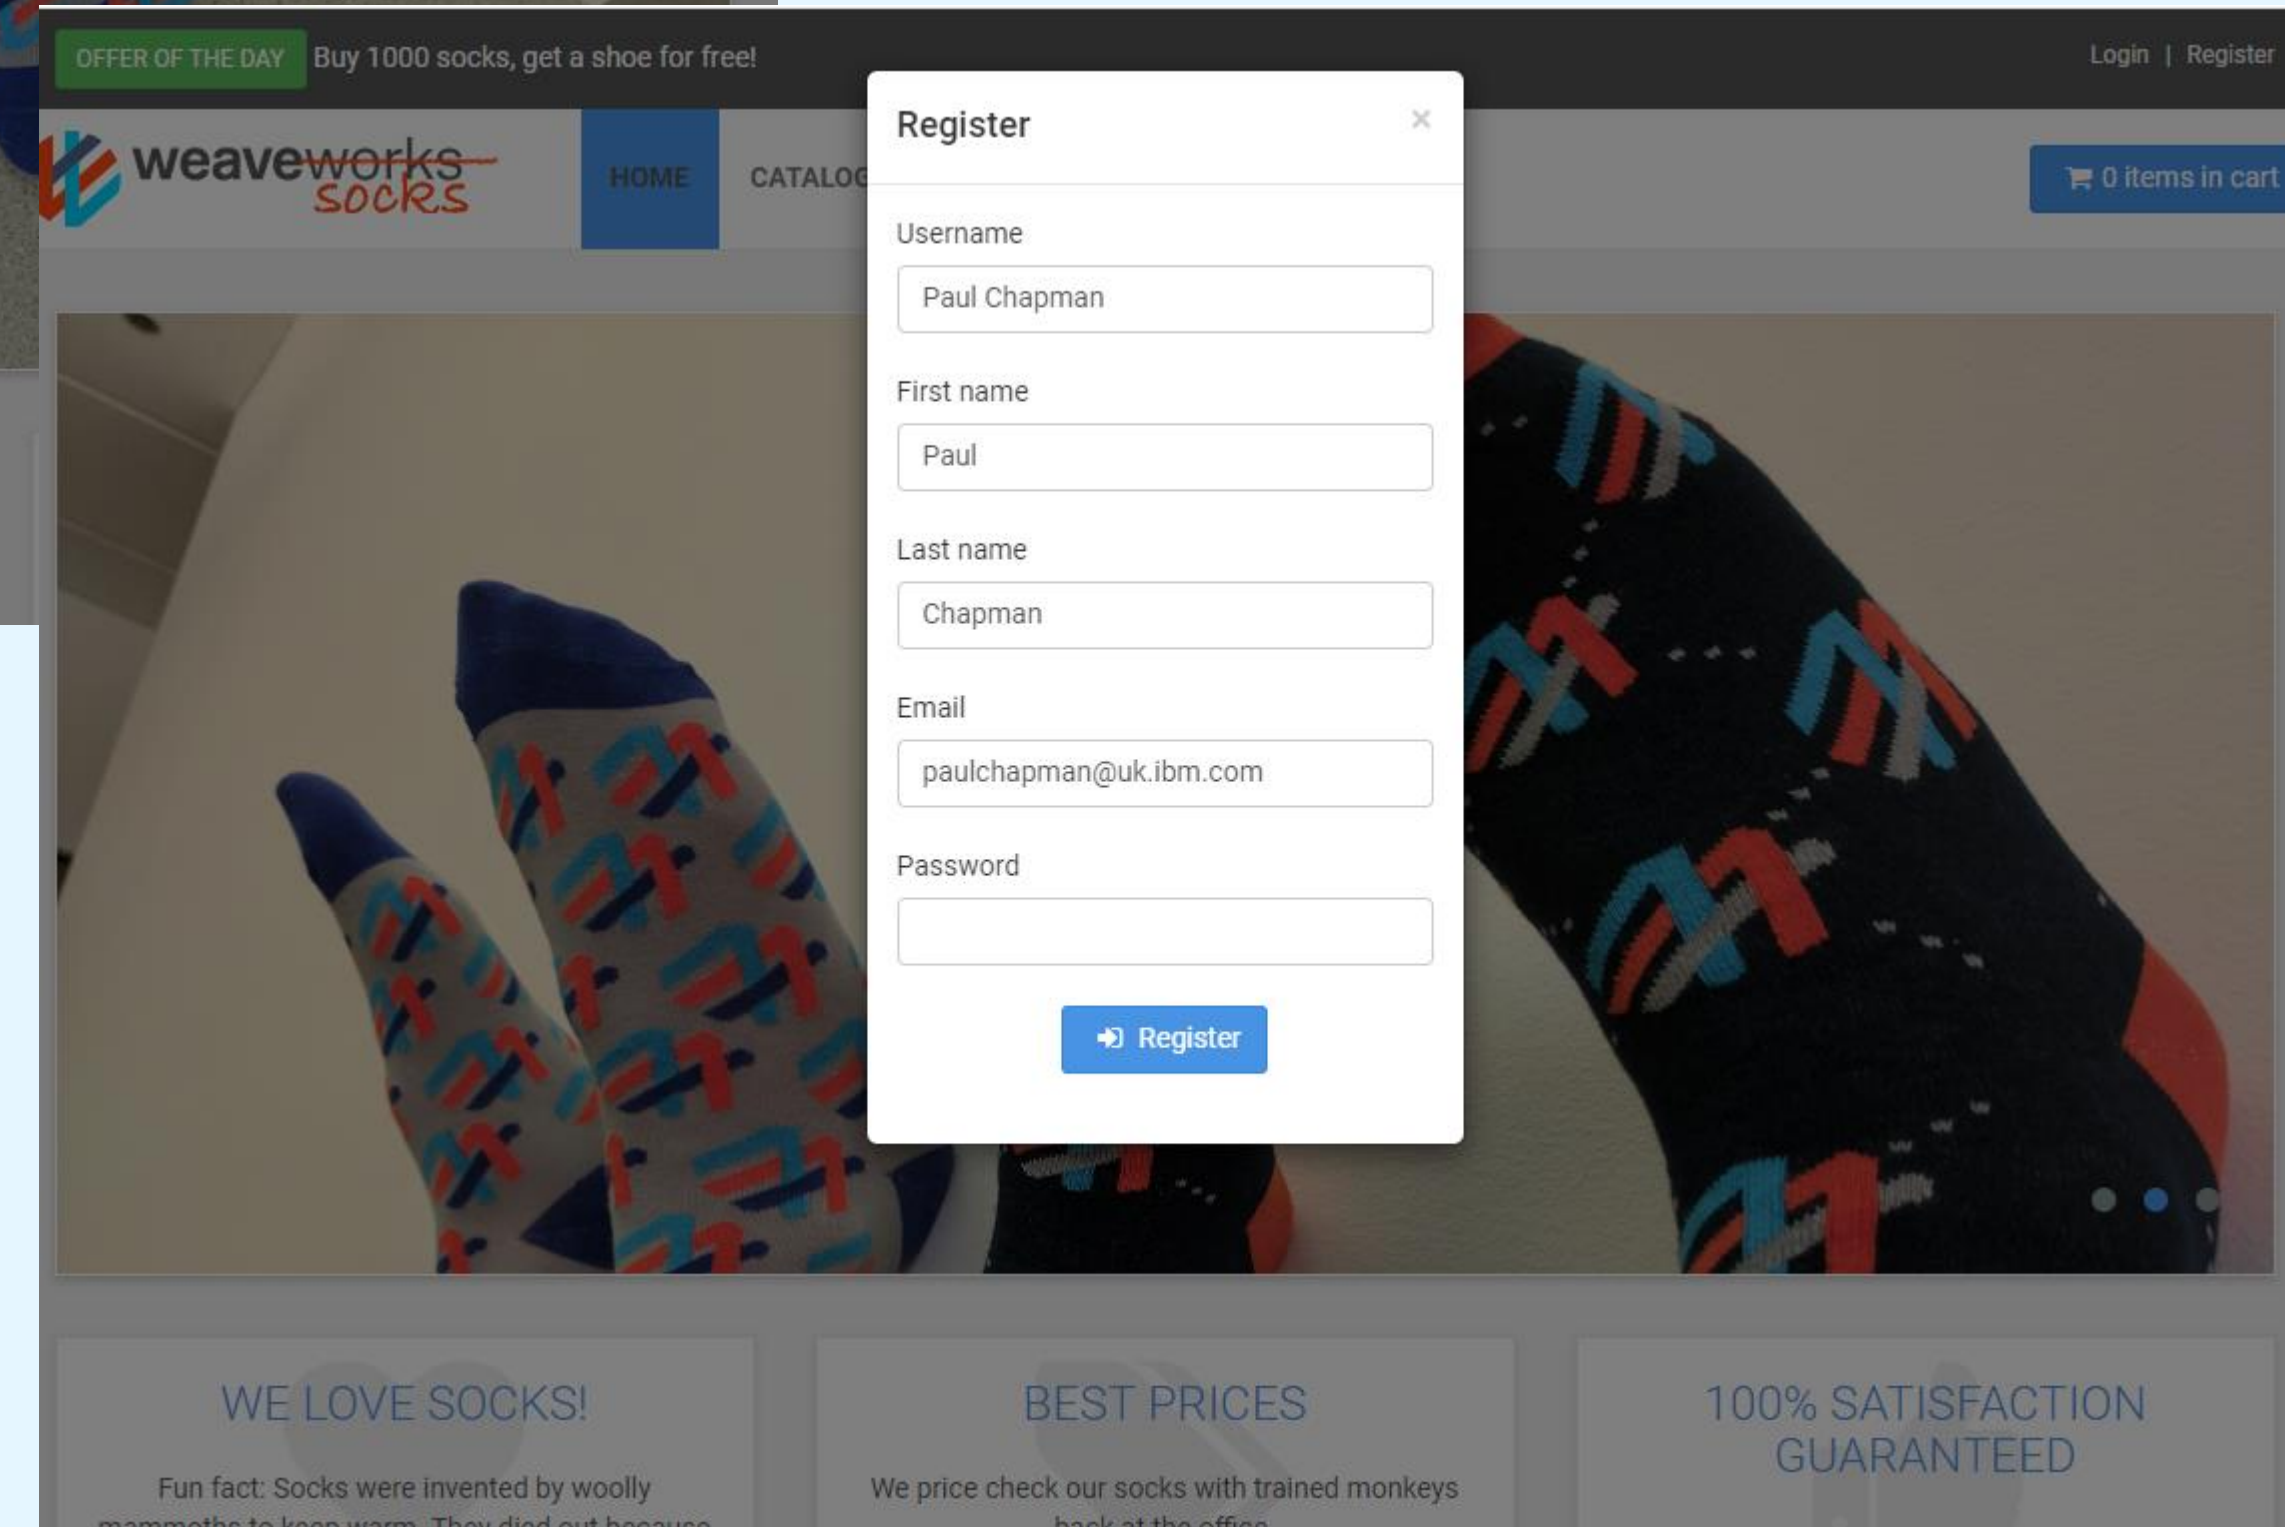
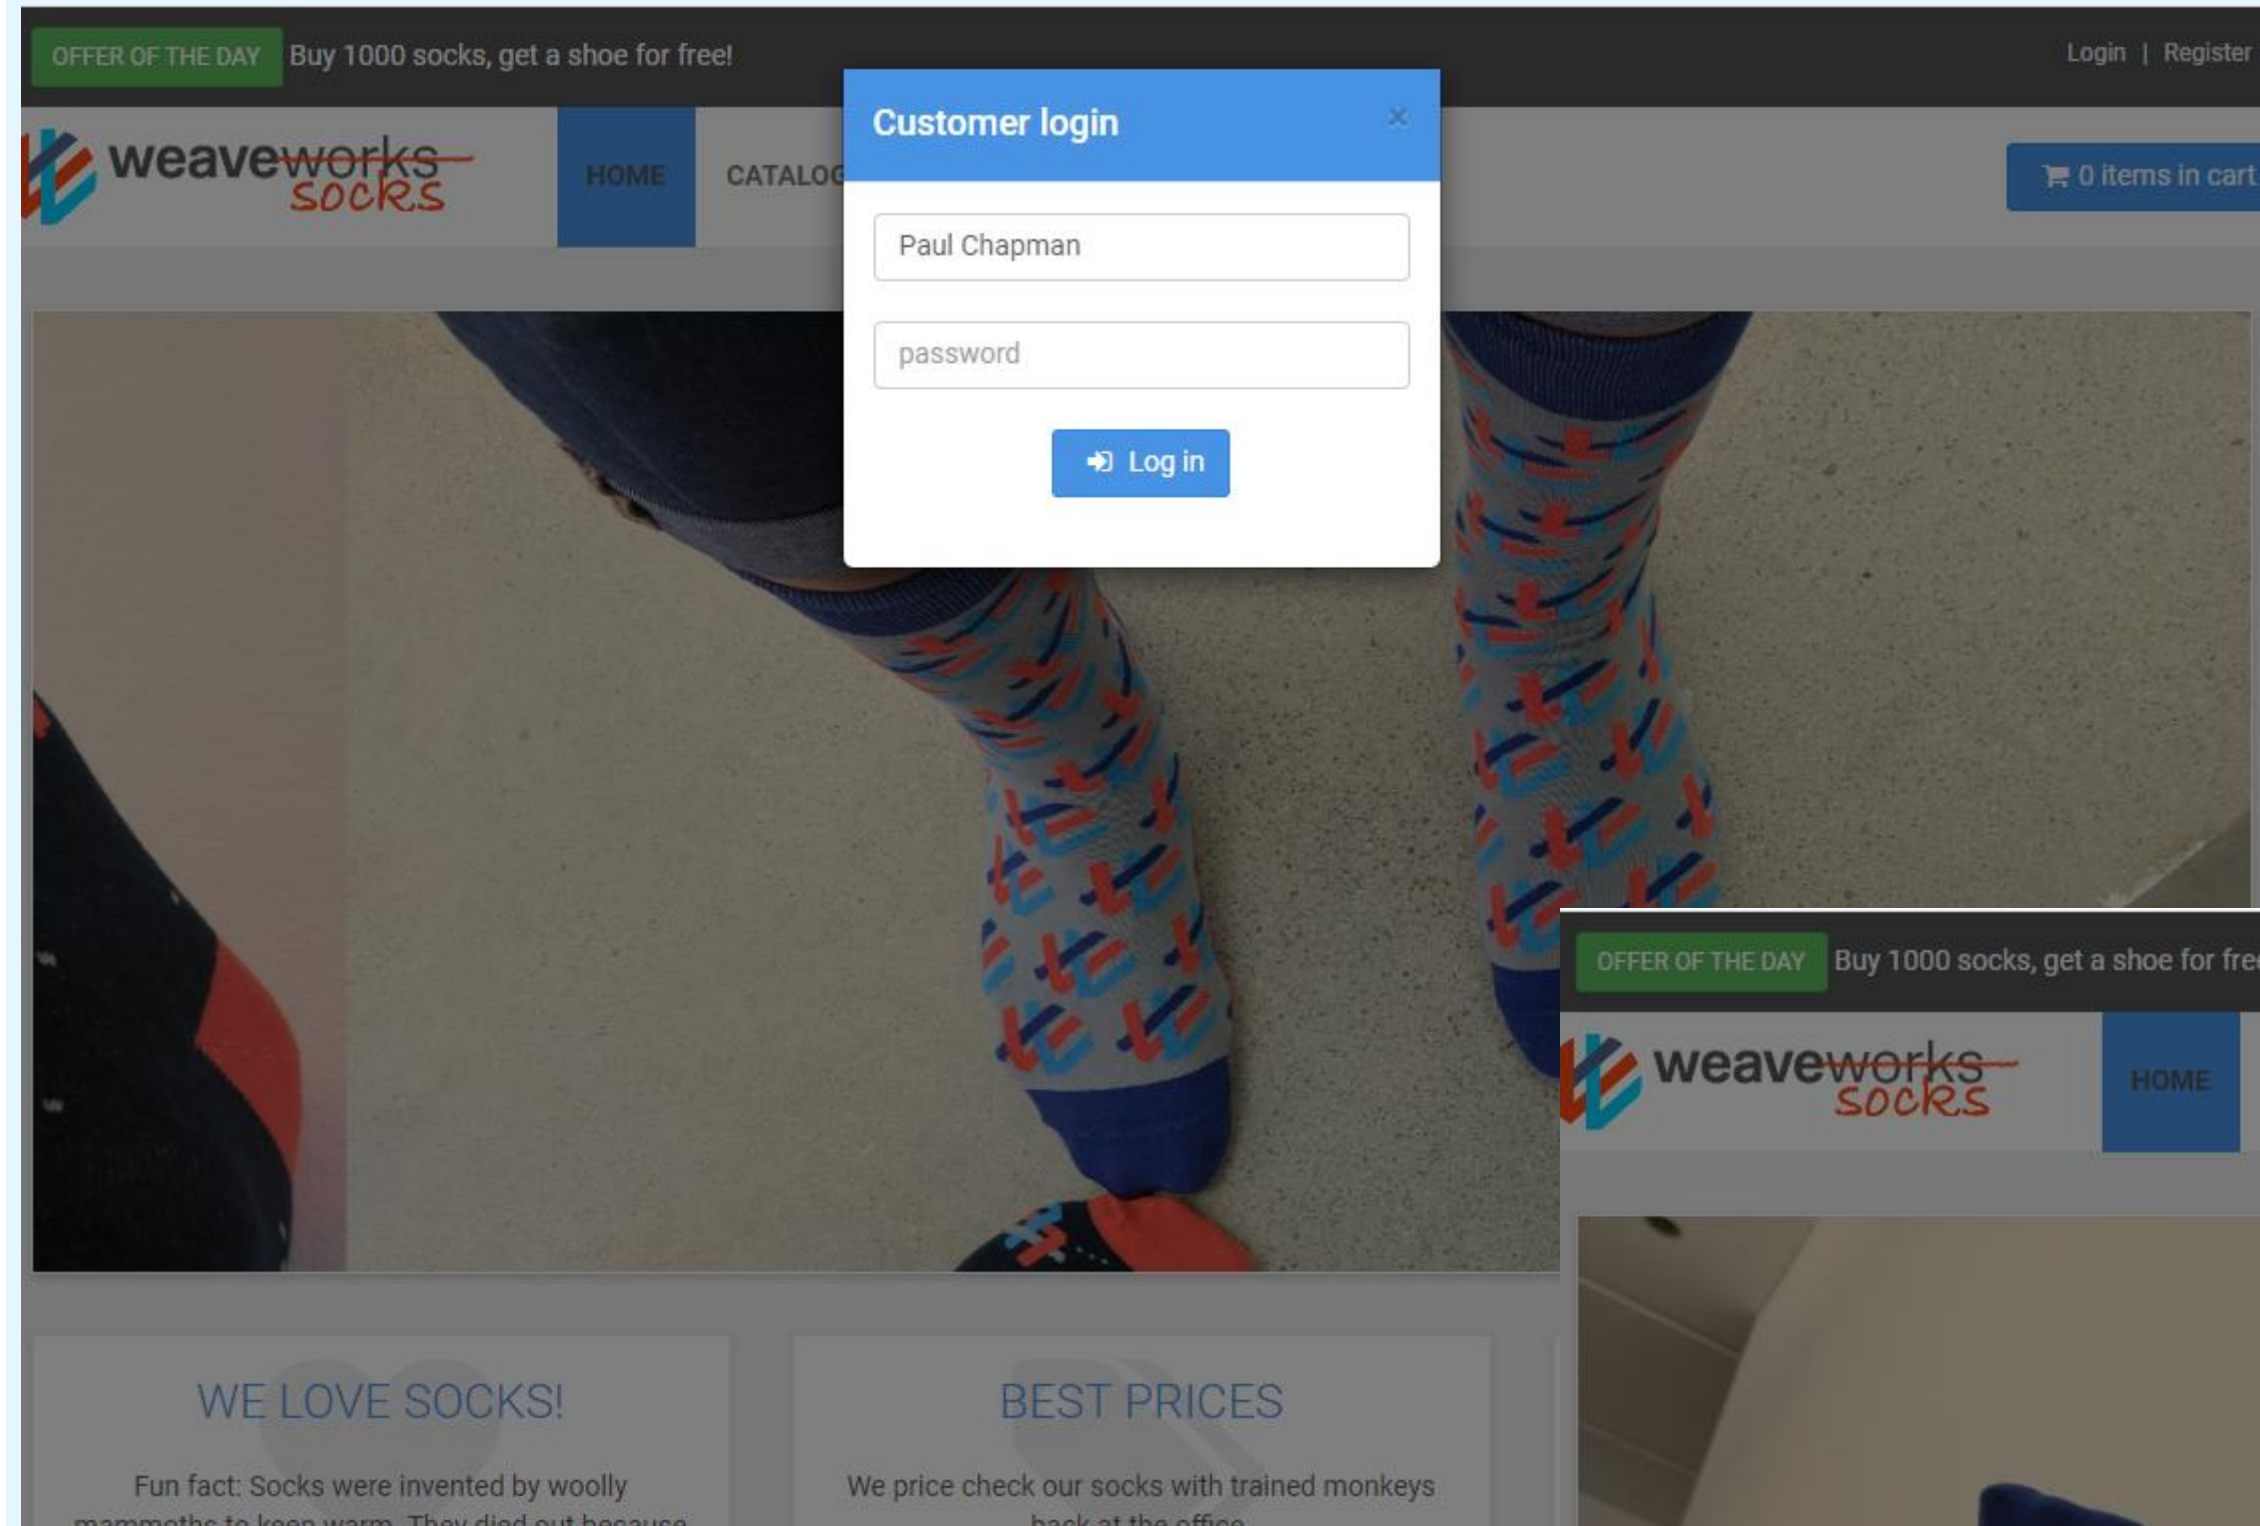
Access Application

- You can access the application by selecting the open URL icon from the front end microservice
- Or you can use CLI to open the URL and paste into your browser



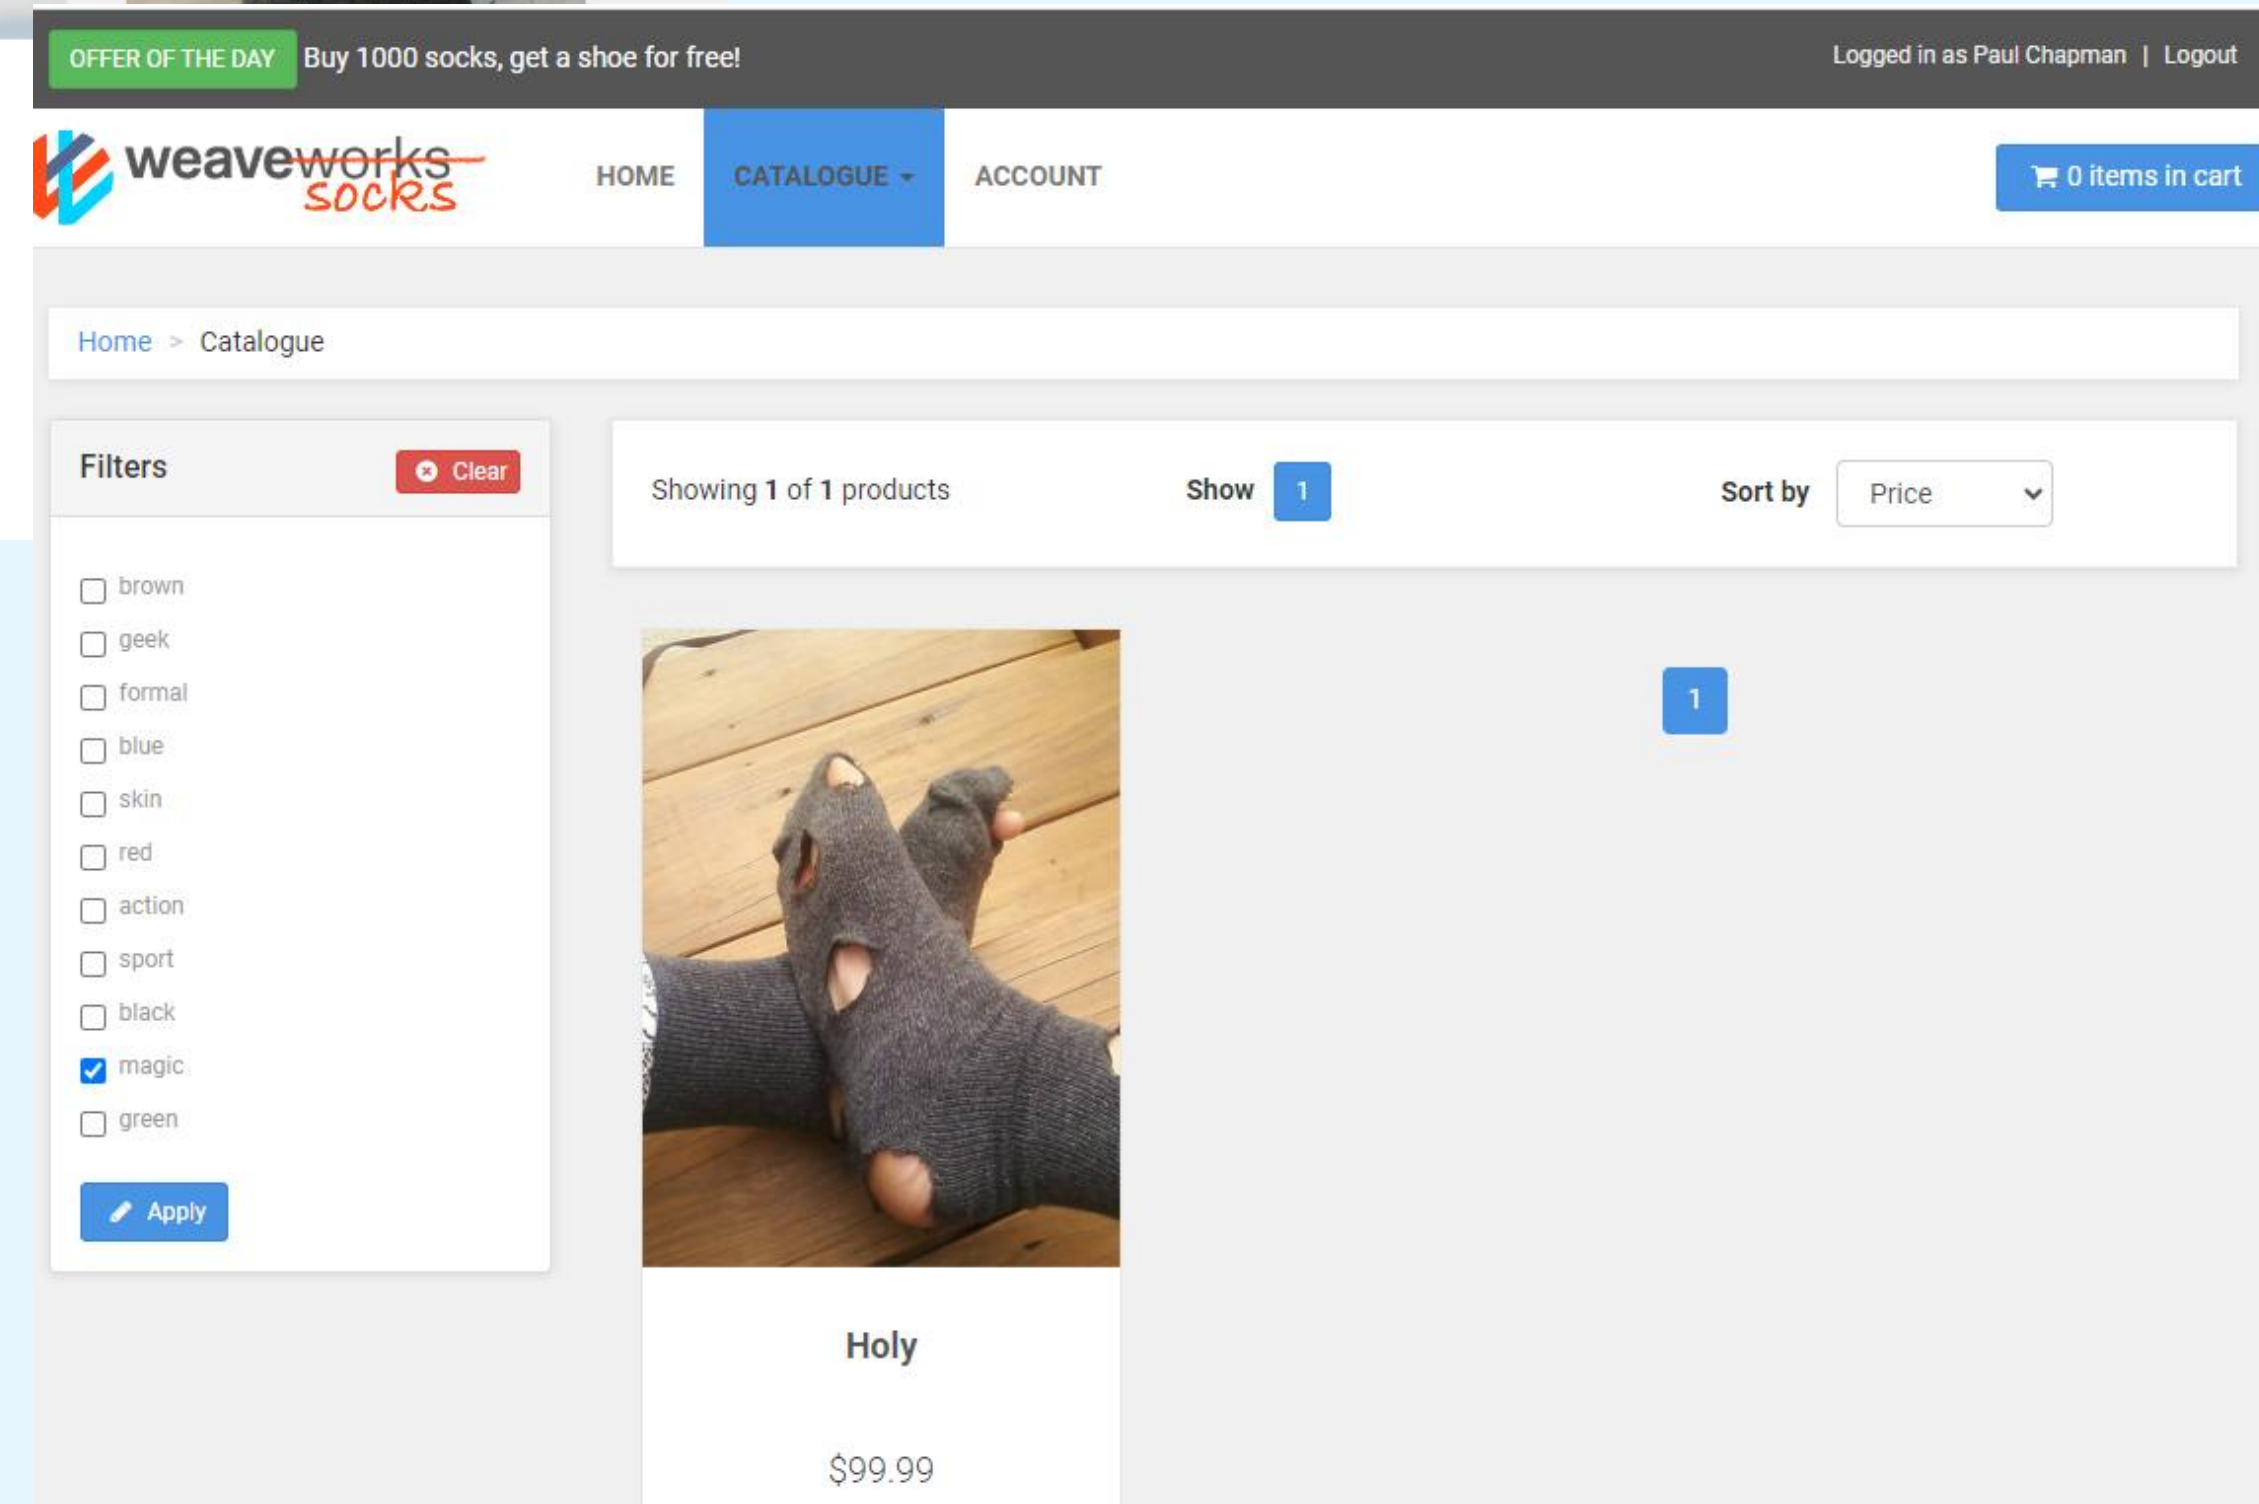
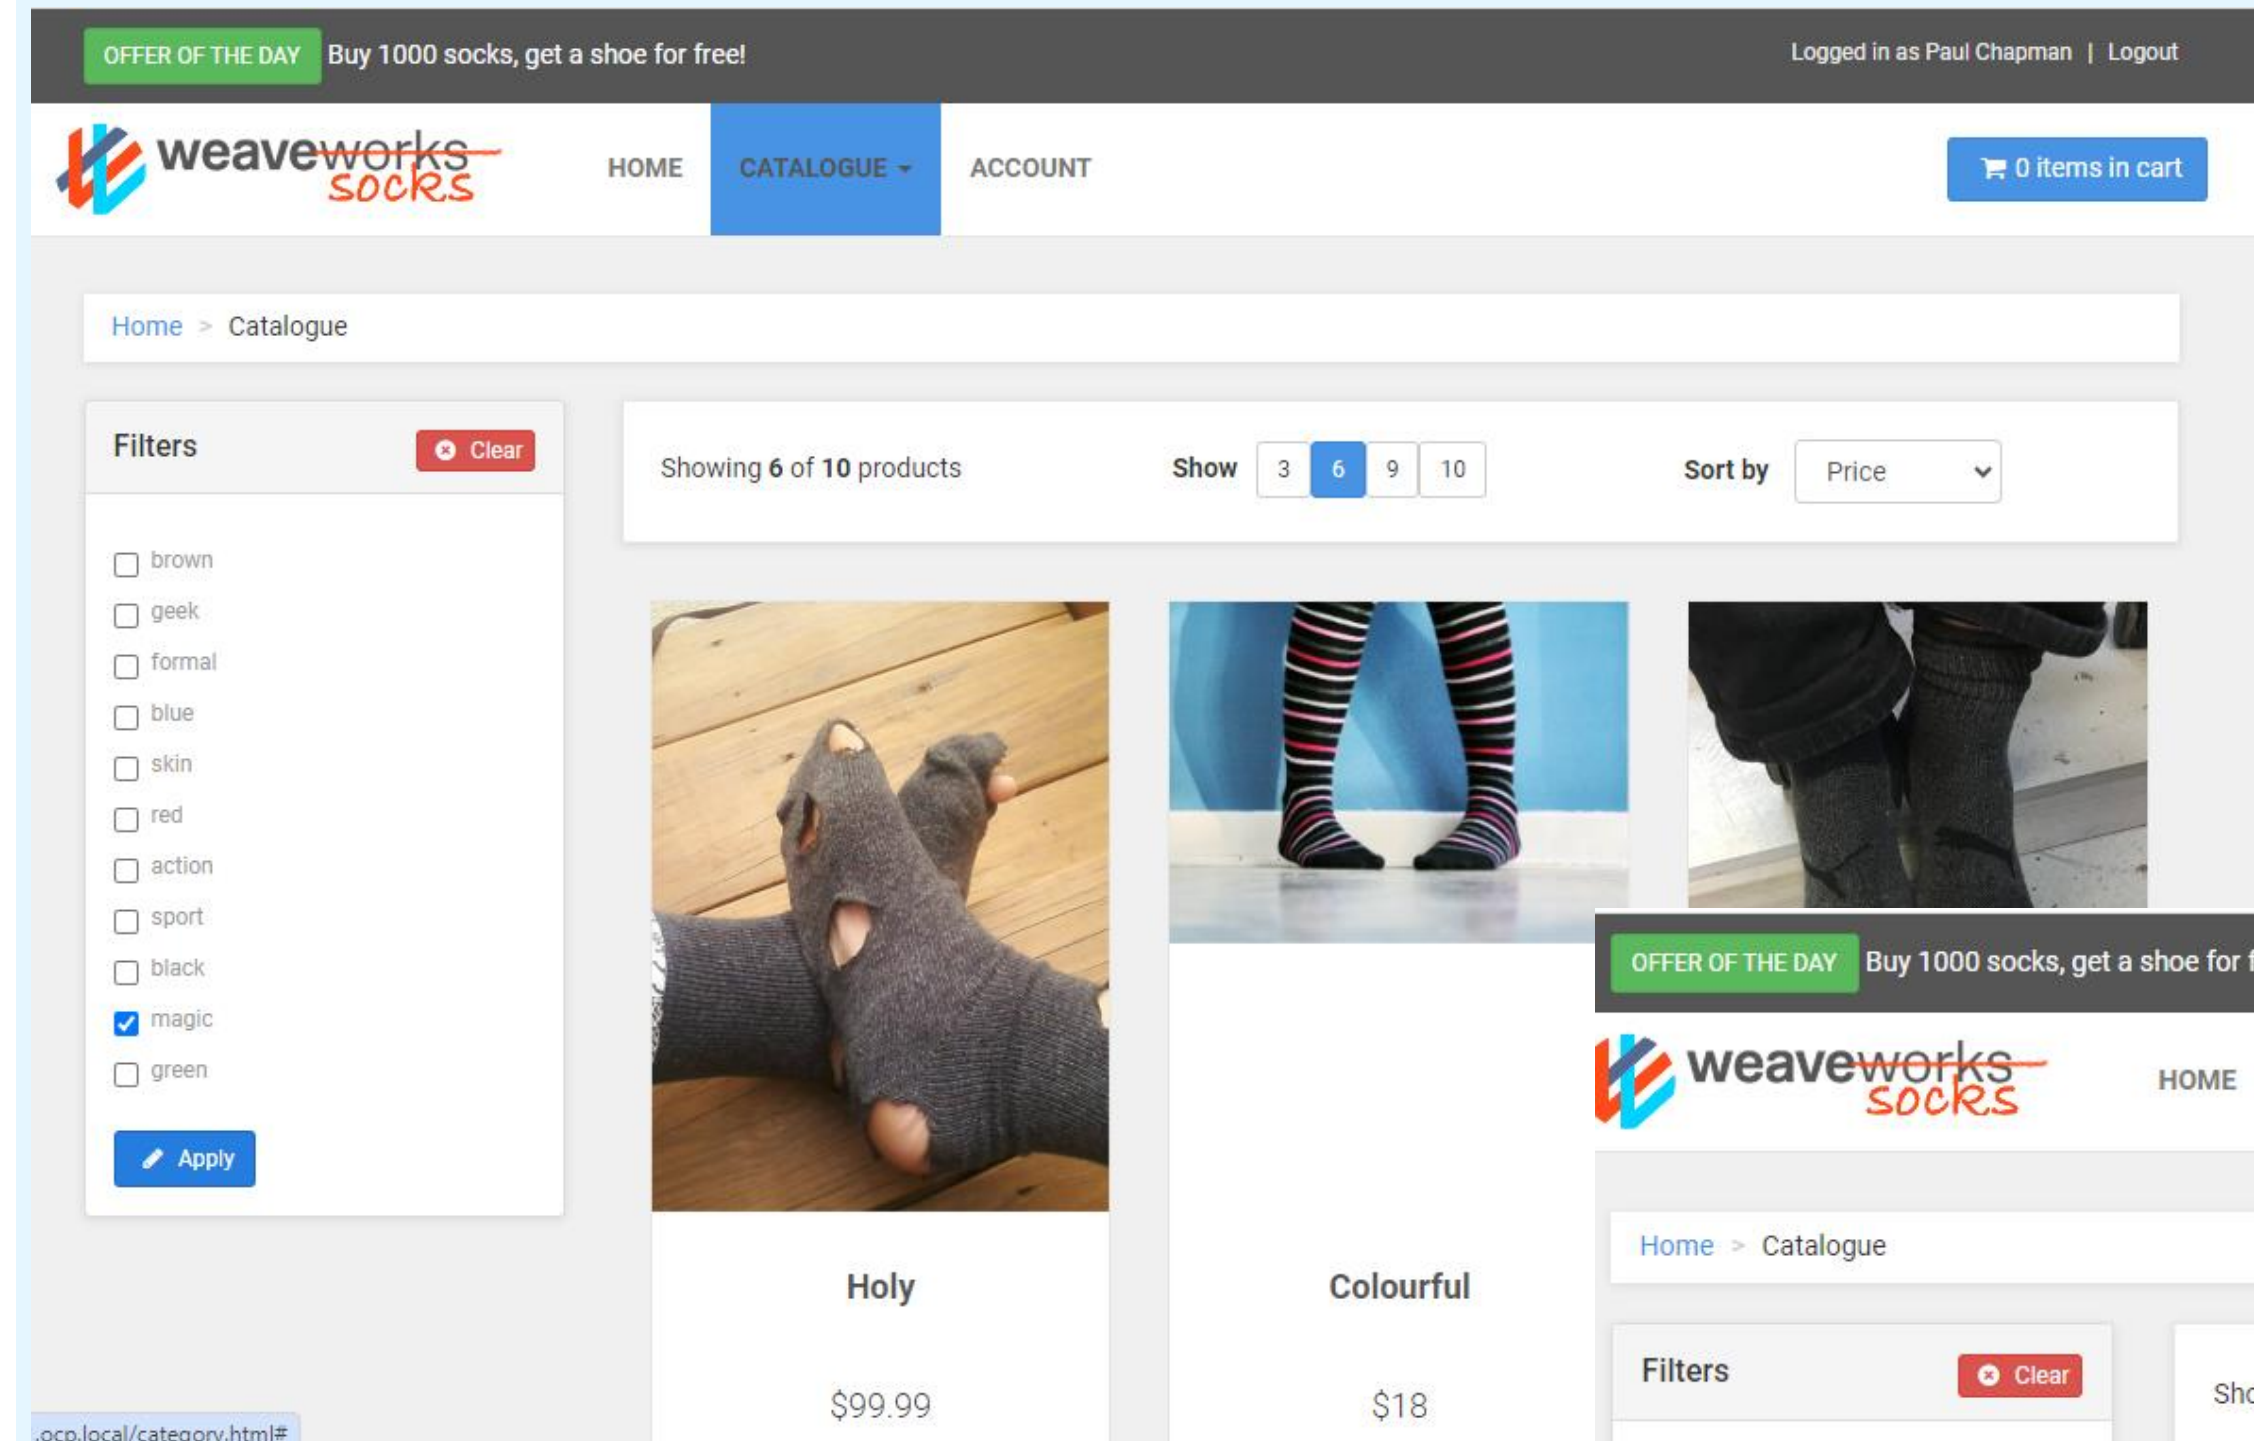
Login or Register

- Login or register as desired
- You will see your username in the top right of the screen once logged in



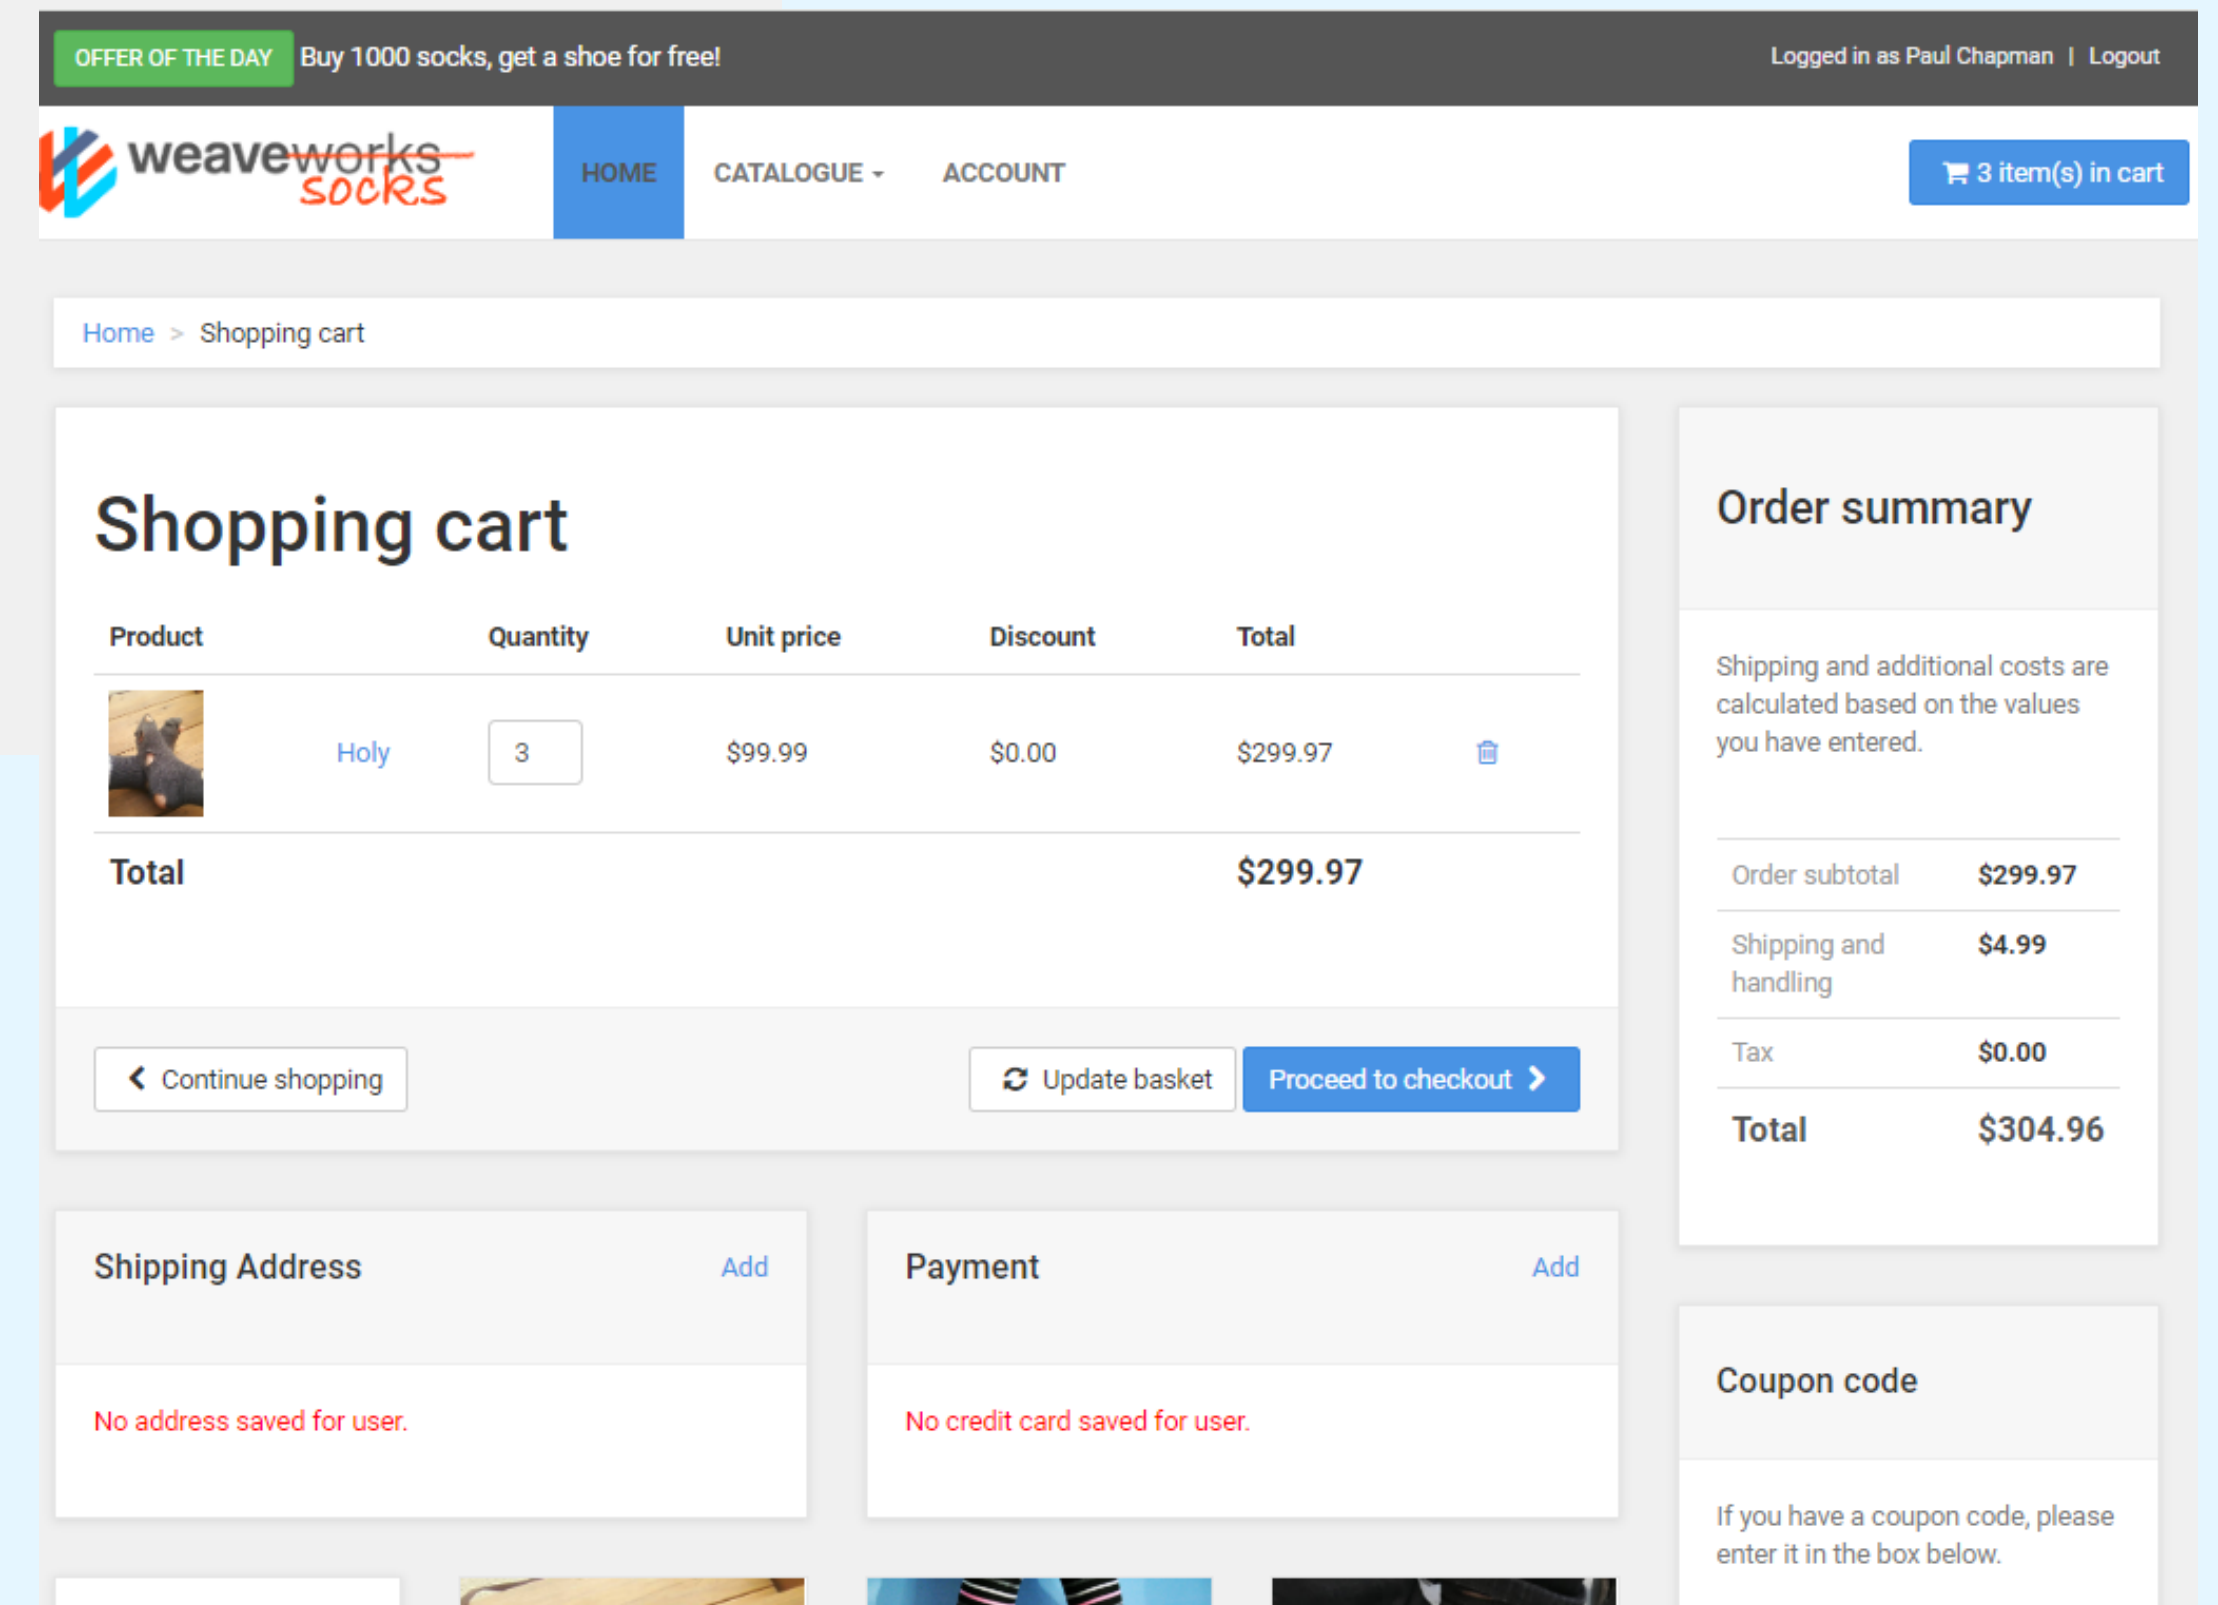
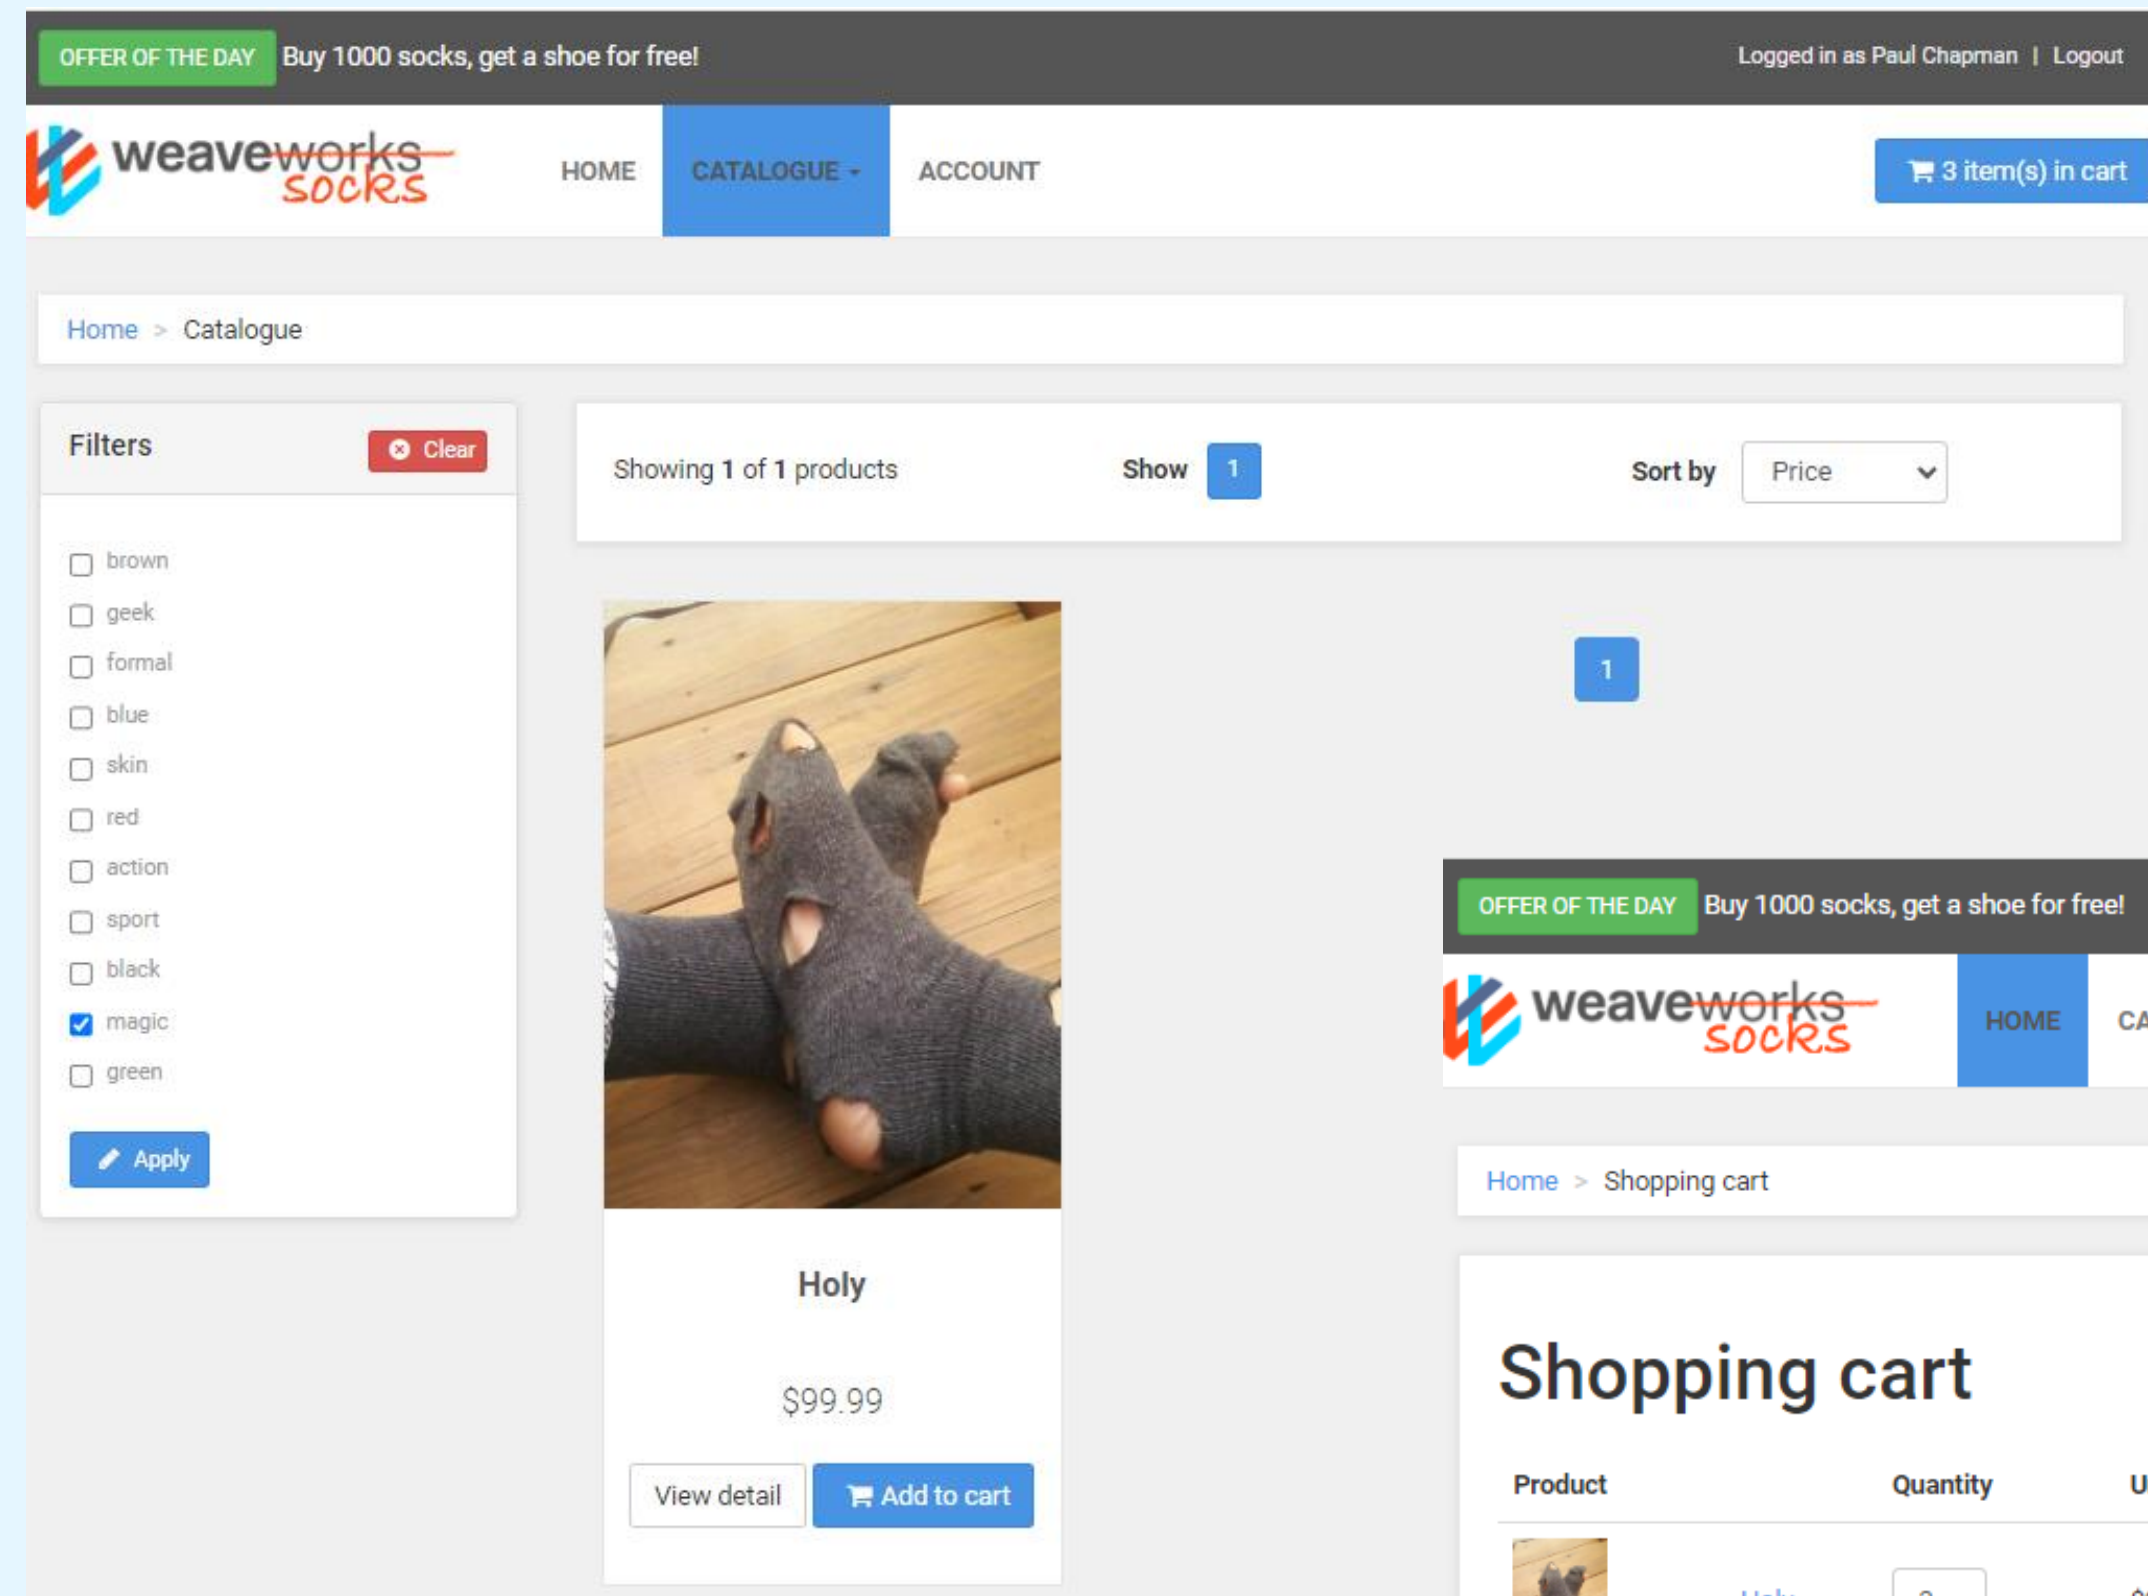
Search Catalogue

- Select options and apply
- Holy socks are very popular, but expensive at \$99



Add to and View Cart

- Select Add to cart
- Note that three items have been added here
- Select Cart to view all item details



Add Payment

- Add details
- Update

The screenshot shows a shopping cart page for 'weaveworks socks'. A modal window titled 'Credit Card' is open, allowing the user to enter payment details. The modal contains three input fields: 'Card Number' (9999888877776666), 'Expires' (10/27), and 'CCV' (111). A blue 'Update' button is located at the bottom right of the modal. The background shows the shopping cart with one item: 'Holy' (quantity 3, unit price \$99.99). The total is \$299.97. The order summary on the right shows a subtotal of \$299.97, shipping and handling of \$4.99, and a total of \$304.96. The shipping address is 1 Old Kent Road, London, W1, UK. The payment section shows 'No credit card saved for user.' and a 'Add' button. The coupon code section is empty.

OFFER OF THE DAY Buy 1000 socks, get a shoe for free!

Logged in as Paul Chapman | Logout


weaveworks socks

HOME CATALOG

3 item(s) in cart

Home > Shopping cart

Shopping cart

Product	Quantity	Unit price
 Holy	3	\$99.99

Total \$299.97

[Continue shopping](#) [Update basket](#) [Proceed to checkout](#)

Shipping Address

1 Old Kent Road
London
W1
UK

Payment

[Add](#)

No credit card saved for user.

Order summary

Shipping and additional costs are calculated based on the values you have entered.

Order subtotal	\$299.97
Shipping and handling	\$4.99
Tax	\$0.00
Total	\$304.96


Coupon code

If you have a coupon code, please enter it in the box below.

Add Address

- Add details
- Update

The screenshot shows a web application interface for 'weaveworks socks'. At the top, there is a promotional banner: 'OFFER OF THE DAY Buy 1000 socks, get a shoe for free!'. The user is logged in as 'Paul Chapman' and has 3 items in their cart. The main content area is titled 'Shopping cart' and contains a table with the following data:

Product	Quantity	Unit price
 Holy	3	\$99.99
Total		

Below the table are buttons for 'Continue shopping' and 'checkout'. An 'Address' modal is open, containing the following form fields:

- House Number: 1
- Street Name: Old Kent Road
- City: London
- Post Code: W1
- Country: UK

An 'Update' button is located at the bottom of the modal. To the right of the modal is an 'Order summary' section:

Order subtotal	\$299.97
Shipping and handling	\$4.99
Tax	\$0.00
Total	\$304.96

Below the order summary are sections for 'Shipping Address' (with an 'Add' button and the message 'No address saved for user.'), 'Payment' (with an 'Add' button and the message 'No credit card saved for user.'), and 'Coupon code' (with a text input field and the instruction 'If you have a coupon code, please enter it in the box below.').

Proceed to Checkout

- Note that order was declined as invoice too high
- Removed two items and updated the cart
- Proceeded to checkout, which was successful this time

The screenshots illustrate the checkout process on the Weaveworks Socks website. The first screenshot shows an error message: "Error placing order. Payment declined: amount exceeds 105.00". The second screenshot shows the shopping cart with one item, "Holy", and a total of \$99.99. The third screenshot shows the "My orders" page with a table of orders, including one that is "Shipped" for \$104.98. The fourth screenshot shows the "My orders" page with a table of orders, including one that is "Shipped" for \$104.98.

Order	Date	Total	Status	Action
# 665c49d6db1e16000176ab5f	2024-06-02 10:30:46	\$ 104.98	Shipped	View

View Order

- View order
- and invoice

OFFER OF THE DAY Buy 1000 socks, get a shoe for free! Logged in as Paul Chapman | Logout

weaveworks socks HOME CATALOGUE - ACCOUNT 0 items in cart

Home > My orders

Customer section

My orders

My orders

Your orders in one place.

If you have any questions, please feel free to [contact us](#), our customer service center is working for you 24/7.

Order	Date	Total	Status	Action
# 665c49d6db1e16000176ab5f	2024-06-02 10:30:46	\$ 104.98	Shipped	View

OFFER OF THE DAY Buy 1000 socks, get a shoe for free! Logged in as Paul Chapman | Logout

weaveworks socks HOME CATALOGUE - ACCOUNT 0 items in cart

Home > My orders > Order


Customer section

My orders

Order #665c49d6db1e16000176ab5f

Order was placed on **2024-06-02 10:30:46** and is currently **shipped**.

If you have any questions, please feel free to [contact us](#), our customer service center is working for you 24/7.

Product	Quantity	Unit price	Discount	Total
 Holy	1	\$99.99	\$0.00	\$99.99
Order subtotal				\$99.99
Shipping and handling				\$4.99
Total				\$104.98

Invoice address	Shipping address
Paul Chapman 1 Old Kent Road London UK	Paul Chapman 1 Old Kent Road London UK

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 [Review Nodes & Pods](#)
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

Review Nodes

- oc get nodes
- Shows type of Node

- oc get nodes -o wide
- shows which arch
 - ppc64le
 - or x86_64

```
paulc:sock-shop-demo$ oc get nodes
NAME          STATUS          ROLES          AGE    VERSION
master-0      Ready          control-plane,master 130d   v1.28.3+4cbdd29
master-1      Ready          control-plane,master 130d   v1.28.3+4cbdd29
master-2      Ready          control-plane,master 130d   v1.28.3+4cbdd29
worker-0      Ready,SchedulingDisabled worker          130d   v1.28.3+4cbdd29
worker-a-0    Ready          worker          94d    v1.28.3+4cbdd29
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc get nodes -o wide
NAME          STATUS          ROLES          AGE    VERSION          INTERNAL-IP          E
XTERNAL-IP    OS-IMAGE          KERNEL-VERSION
CONTAINER-RUNTIME
master-0      Ready          control-plane,master 130d   v1.28.3+4cbdd29  10.20.181.156  <
none>        Red Hat Enterprise Linux CoreOS 415.92.202311061558-0 (Plo
w)          5.14.0-284.40.1.el9_2.ppc64le
cri-o://1.28.2-2.rhaos4.15.gite7be4e1.el9
master-1      Ready          control-plane,master 130d   v1.28.3+4cbdd29  10.20.176.68   <
none>        Red Hat Enterprise Linux CoreOS 415.92.202311061558-0 (Plo
w)          5.14.0-284.40.1.el9_2.ppc64le
cri-o://1.28.2-2.rhaos4.15.gite7be4e1.el9
master-2      Ready          control-plane,master 130d   v1.28.3+4cbdd29  10.20.176.132  <
none>        Red Hat Enterprise Linux CoreOS 415.92.202311061558-0 (Plo
w)          5.14.0-284.40.1.el9_2.ppc64le
cri-o://1.28.2-2.rhaos4.15.gite7be4e1.el9
worker-0      Ready,SchedulingDisabled worker          130d   v1.28.3+4cbdd29  10.20.176.243  <
none>        Red Hat Enterprise Linux CoreOS 415.92.202311061558-0 (Plo
w)          5.14.0-284.40.1.el9_2.ppc64le
cri-o://1.28.2-2.rhaos4.15.gite7be4e1.el9
worker-a-0    Ready          worker          94d    v1.28.3+4cbdd29  10.20.29.114   <
none>        Red Hat Enterprise Linux CoreOS 415.92.202311061558-0 (Plo
w)          5.14.0-284.40.1.el9_2.x86_64
cri-o://1.28.2-2.rhaos4.15.gite7be4e1.el9
paulc:sock-shop-demo$
```

Review Pods

- `oc get pods`
- You can also specify to show Pods running on particular Worker Nodes
- In this case;
 - worker-0 is ppc64le
 - worker-a-0 is x86
- Or specify a pod to determine which Node (and architecture it's running on

```
paulc:sock-shop-demo$ oc get pods --field-selector=spec.host=worker-0
```

NAME	READY	STATUS	RESTARTS	AGE
carts-ddfd7779f-4t2pm	1/1	Running	0	5h37m
payment-5bb5685f8b-hwhzv	1/1	Running	0	5h37m
queue-master-79c7ffcdbe-cs2fj	1/1	Running	0	5h37m
shipping-5d59b847c6-rc6gk	1/1	Running	0	5h37m

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$ oc get pods --field-selector=spec.host=worker-a-0
```

NAME	READY	STATUS	RESTARTS	AGE
carts-db-845dc96575-z9ngc	1/1	Running	0	5h38m
catalogue-7db4b66745-vct85	1/1	Running	0	5h38m
catalogue-db-d6977946d-qf2wb	1/1	Running	0	5h38m
front-end-6565bc7c48-zk46h	1/1	Running	0	5h38m
orders-7b898d9d8b-2g9nv	1/1	Running	0	5h38m
orders-db-7d9d78bb9f-pc7nv	1/1	Running	0	5h38m
rabbitmq-6d9b8c7d9-nvqk8	2/2	Running	0	5h37m
session-db-6bf7c5779f-v6t4f	1/1	Running	0	5h13m
user-86df5b8df4-qsznq	1/1	Running	0	5h37m
user-db-6fbdd4cb65-zmdtk	1/1	Running	0	5h37m

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$
```

```
paulc:sock-shop-demo$ oc get pods -o wide -l name=front-end
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	NOMINATED NODE
front-end-6565bc7c48-zk46h	1/1	Running	0	5h43m	10.129.3.118	worker-a-0	<none>

```
paulc:sock-shop-demo$
```

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 [Cordon, Migration Between Architectures](#)
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 Additional Resources

Cordon Nodes

- Front End is hosted by x86 worker-a-0
- Cordon x86 Node
- Scheduling is now disabled on x86

```
paulc:sock-shop-demo$ oc get pods -o wide -l name=front-end
NAME                                READY   STATUS    RESTARTS   AGE   IP             NODE           NOMINATED NODE
E   READINESS GATES
front-end-6565bc7c48-xmhb6          1/1     Running   0           54s   10.129.3.175   worker-a-0     <none>
<none>
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc adm cordon worker-a-0
node/worker-a-0 cordoned
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc get nodes
NAME                STATUS              ROLES                    AGE   VERSION
master-0            Ready               control-plane,master    130d  v1.28.3+4cbdd29
master-1            Ready               control-plane,master    130d  v1.28.3+4cbdd29
master-2            Ready               control-plane,master    130d  v1.28.3+4cbdd29
worker-0            Ready               worker                   130d  v1.28.3+4cbdd29
worker-a-0          Ready,SchedulingDisabled worker                   94d   v1.28.3+4cbdd29
paulc:sock-shop-demo$
```


Move from x86 to Power

- Front End is hosted by x86 worker-a-0
- Delete Front End, forcing restart on Power
- Front End is now running on Power Worker Node

```
paulc:sock-shop-demo$ oc get pods -o wide -l name=front-end
NAME                                READY   STATUS    RESTARTS   AGE       IP            NODE           NOMINATED NODE
ODE  READINESS GATES
front-end-6565bc7c48-xmhb6          1/1     Running   0           3m50s    10.129.3.175  worker-a-0    <none>
<none>
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc delete pod front-end-6565bc7c48-xmhb6
pod "front-end-6565bc7c48-xmhb6" deleted
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc get pods -o wide -l name=front-end
NAME                                READY   STATUS    RESTARTS   AGE       IP            NODE           NOMINATED NODE
ODE  READINESS GATES
front-end-6565bc7c48-4pcvp          1/1     Running   0           37s      10.131.0.168  worker-0      <none>
<none>
paulc:sock-shop-demo$
```

Uncordon x86 Node

- Scheduling on x86 is disabled
- Uncordon x86 Node
- X86 Nodes is now schedulable again

```
paulc:sock-shop-demo$ oc get nodes
NAME                STATUS              ROLES                    AGE      VERSION
master-0            Ready              control-plane,master    130d    v1.28.3+4cbdd29
master-1            Ready              control-plane,master    130d    v1.28.3+4cbdd29
master-2            Ready              control-plane,master    130d    v1.28.3+4cbdd29
worker-0            Ready              worker                   130d    v1.28.3+4cbdd29
worker-a-0          Ready,SchedulingDisabled  worker                   94d     v1.28.3+4cbdd29
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc adm uncordon worker-a-0
node/worker-a-0 uncordoned
paulc:sock-shop-demo$
paulc:sock-shop-demo$
paulc:sock-shop-demo$ oc get nodes
NAME                STATUS              ROLES                    AGE      VERSION
master-0            Ready              control-plane,master    130d    v1.28.3+4cbdd29
master-1            Ready              control-plane,master    130d    v1.28.3+4cbdd29
master-2            Ready              control-plane,master    130d    v1.28.3+4cbdd29
worker-0            Ready              worker                   130d    v1.28.3+4cbdd29
worker-a-0          Ready              worker                   94d     v1.28.3+4cbdd29
paulc:sock-shop-demo$
```

Use: Node Selectors

Add the `nodeSelector` field and add an architecture limitation using a `Node` label:

nodeSelector:

```
node.openshift.io/os_id: rhcos  
kubernetes.io/arch: amd64
```

nodeSelector:

```
node.openshift.io/os_id: rhcos  
kubernetes.io/arch: ppc64le
```

For each Pod... it directs the workload to the node that matches the architecture.

<https://community.ibm.com/community/user/powerdeveloper/blogs/paul-bastide/2024/01/09/multi-arch-compute-node-selector>

Use: Taints and Tolerations

[Node affinity](#) is a property of [Pods](#) that *attracts* them to a set of [nodes](#) (either as a preference or a hard requirement). *Taints* are the opposite -- they allow a node to repel a set of pods.


Tolerations are applied to pods. Tolerations allow the scheduler to schedule pods with matching taints. Tolerations allow scheduling but don't guarantee scheduling: the scheduler also [evaluates other parameters](#) as part of its function.

Taints and tolerations work together to ensure that pods are not scheduled onto inappropriate nodes. One or more taints are applied to a node; this marks that the node should not accept any pods that do not tolerate the taints.

```
oc adm taint nodes worker-amd-0 kubernetes.io/arch=ppc64le:NoSchedule
```

This means that no pod will be able to schedule onto node1 unless it has a matching toleration.

Multi-Arch Tuning Operator



Multiarch Tuning Operator

0.9.0 provided by Red Hat

[Install](#)

Channel

tech-preview

Version

0.9.0

Capability level

- Basic Install
- Seamless Upgrades
- Full Lifecycle
- Deep Insights
- Auto Pilot

Source

Multiarch Tuning Operator
Catalog

The Multiarch Tuning Operator enhances the user experience for administrators of Openshift clusters with multi-architecture compute nodes or Site Reliability Engineers willing to migrate from single-arch to multi-arch OpenShift. When diverse CPU architectures coexist within a cluster, the Multiarch Tuning Operator stands out as a pivotal tool to enhance efficiency and streamline operations such as architecture-aware scheduling of workloads.

Operands

- Architecture aware Pod Placement:** The pod placement operand consists of the `PodPlacementController` and the `PodPlacementWebhook` and is managed through a singleton custom resource - `podplacementconfigs.multiarch.openshift.io`. Its aim is to automate the set up of strong predicates based on the `kubernetes.io/arch` label in the pod's `nodeAffinity` by inspecting the container images in each pod and deriving a set of architectures supported by the pod. When a pod is created, the `PodPlacementWebhook` will add the `multiarch.openshift.io/scheduling-gate` scheduling gate. It will prevent the pod from being scheduled until the `PodPlacementController` computes a predicate for the `kubernetes.io/arch` label, adds it as a node affinity requirement to the pod spec, and removes the scheduling gate.

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 [Multi-Arch Compute Installation](#)
- 09 Public Reference
- 10 Additional Resources

Installation

Documentation

The screenshot shows the Red Hat OpenShift documentation page for "Creating a cluster with multi-architecture compute machines on IBM Power". The page includes a navigation menu on the left with categories like "About", "Release notes", "Getting started", "Architecture", "Installing", and "Post-installation configuration". The main content area features a search bar, a breadcrumb trail, and a table of contents. The primary text explains that to create a multi-architecture cluster on IBM Power (ppc64le), an existing single-architecture (x86_64) cluster must be upgraded. An "IMPORTANT" note states that the cluster must be upgraded to use the multi-architecture payload. A "NOTE" at the bottom indicates that hosts for different architectures must share the same storage layer. A "Verifying cluster compatibility" section is also visible.

https://docs.openshift.com/container-platform/4.14/post_installation_configuration/configuring-multi-arch-compute-machines/creating-multi-arch-compute-nodes-ibm-power.html



NOTE

You should limit the number of network hops between the compute and control plane as much as possible.

Procedure

- You can check that your cluster uses the architecture payload by running the following command:

```
$ oc adm release info -o jsonpath="{.metadata.metadata}"
```

Verification

- If you see the following output, then your cluster is using the multi-architecture payload:

```
{
  "release.openshift.io/architecture": "multi",
  "url": "https://access.redhat.com/errata/<errata_version>"
}
```

You can then begin adding multi-arch compute nodes to your cluster.

- If you see the following output, then your cluster is not using the multi-architecture payload:

```
{
  "url": "https://access.redhat.com/errata/<errata_version>"
}
```

Cost & Subscription Considerations

- Multi-Arch Compute is included with OCP
- x86 and ARM share Red Hat SKU
- Power & Z have different SKU's

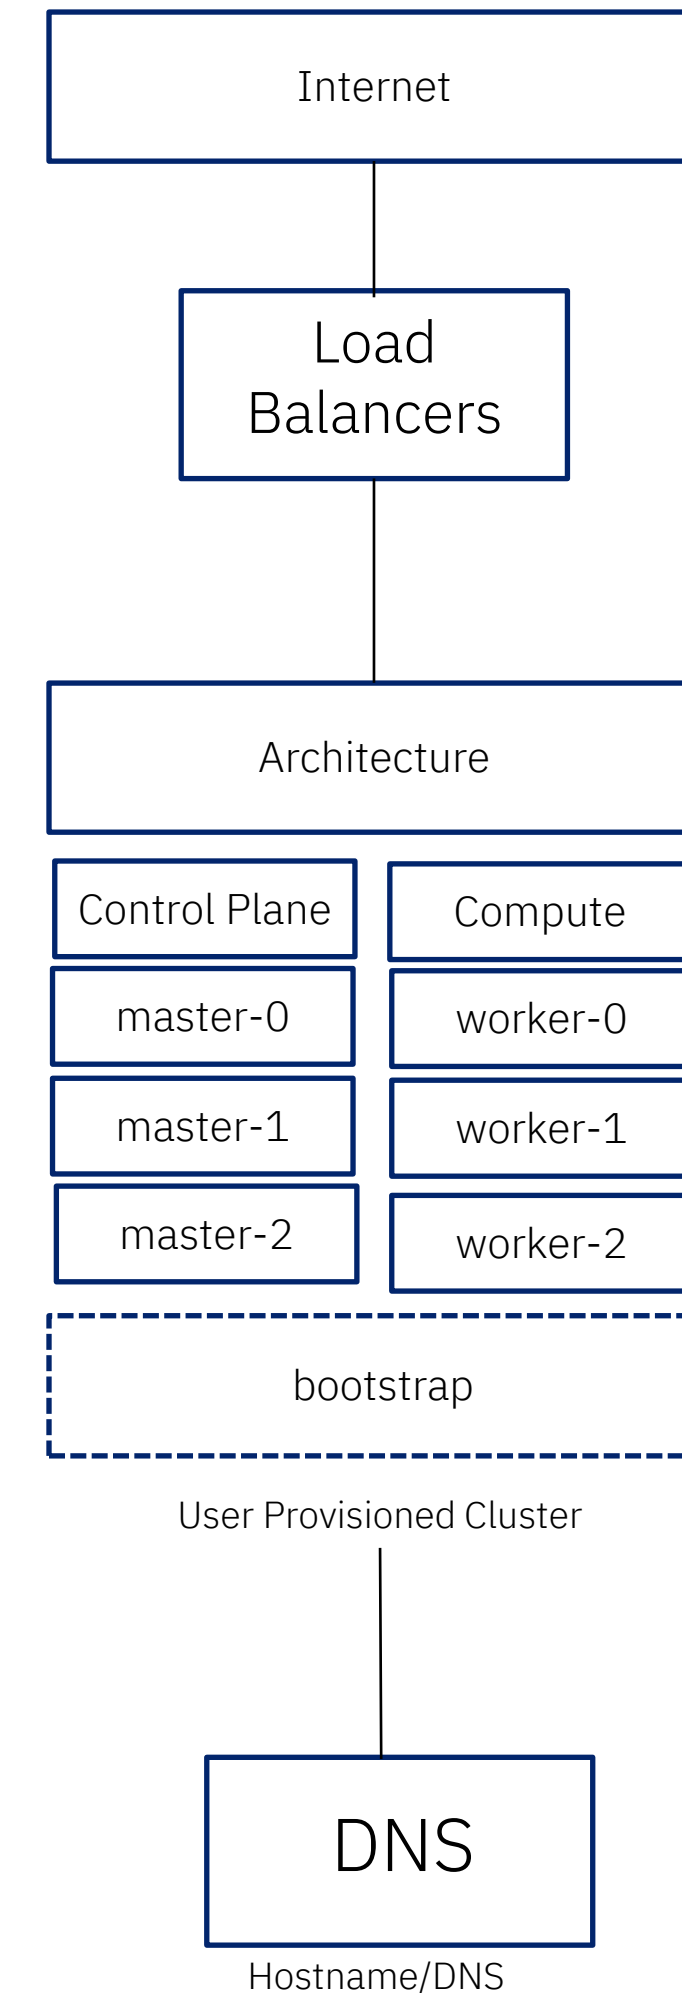
Installation

A simple cluster is installed on a single target architecture.

The cluster has a homogenous control plane and homogenous compute plane.

Front End is load balanced and supported by DNS entries pointing to the Ingress/API/MachineConfig Server.

~45-1 hour post installation



Post Installation Multi-Arch Compute

1. Prepare
 1. Networking
 2. Bastion for Post Installation Support Services
 3. Prepare Ignition
 4. Migrate to the Multi Payload
 5. Prepare Cluster Services
2. Image
 1. Download Architecture specific Image
 2. Load Image in Target Platform
3. Ignite Workers
 1. Start them up
 2. Approve Node Bootstrapper
 3. Issue Kubelet Certificate

Almost always a UPI setup.

With a converged Hyperscaler, single service, you can use MachineSets to scale up and down.

Step 3 in this case is automated.

Start using the new Workers.

~45-1 hour post installation

Setup: Automation

The Power Hybrid Cloud team has added three main use-cases. We continue to refine these repositories to best support our daily CI/QE.

We accept issues/comments/commits on the code as we refine it.

Add Power Node to IBM Cloud VPC IPI

<https://github.com/ibm/ocp4-upi-compute-powervs>

Currently being hardened.

Add Intel Node to IBM PowerVS Workspace UPI

<https://github.com/ibm/ocp4-upi-compute-powervs-ibmcloud>

Add OpenStack Intel Node to IBM PowerVM/PowerVC UPI Cluster

<https://github.com/ocp-power-automation/ocp4-upi-multiarch-compute/tree/main>

Work in Progress

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 [Public Reference](#)
- 10 Additional Resources

Try Multi-Architecture Clusters

Would you like to join the Power MAC Early Adoption Program?

1. Meet with Development
2. Demo
3. 4-week trial environment
4. Provide feedback

Run x86, ARM & Power Containers within the same OpenShift Cluster


- Reduces costs
- Enables complex solutions with some components that do not run on Power, to run on Power
- Enables simpler migration from x86 to Power

Invitation to Participate in OpenShift Multi-Architecture Clusters

Trial, Collaboration & Feedback



Paul Chapman
IBM, Global Power Modernization Technical Lead



– <https://www.youtube.com/watch?v=MVrRJzQAJg8>

Public Reference- smeup

IBM Power


Bringing Multi-Architecture Compute to IBM Power

The release of [Red Hat OpenShift 4.14](#) brought the OpenShift Container Platform Multi-Architecture Compute (MAC) feature to IBM Power. Multi-Architecture Compute provides a single heterogeneous cluster, enabling fit-for-purpose computing so clients can align tasks and applications to CPU strengths and software availability rather than one architecture.

With clients having expressed interest in incorporating the feature into their solution on Power, IBM kicked off an early adoption program to enable [SME.up](#) for Multi-Architecture Compute, obtain feedback on the usability of MAC, receive input on onboarding materials, and further validate and refine use cases. Led by the IBM Power Development and Design teams, IBM Power's Multi-Arch Cluster Early Adoption and Feedback Program allowed clients and partners to conduct a trial and provide feedback. Working with SME.up, an IBM Power partner based in Italy, the team evaluated the interest in mixing architectures and the instances in which clients sought to implement MAC on Power. SME.up has a long history of configuring and deploying solutions on IBM Power, positioning them as the perfect partner for this opportunity to enhance usability moving forward.

Getting started

Working to co-create MAC on Power, the IBM Power Development and Design teams provided SME.up with step-by-step instructions and architecture diagrams to manufacture the onboarding process. SME.up began by first setting up the MAC cluster and deploying an existing application before progressing further and successfully deploying a new application in the MAC cluster on Power.



Engaging in co-creation proved to be mutually beneficial. Through the lens of the IBM teams, SME.up provided invaluable feedback, which led to a more streamlined onboarding documentation. From the partner's perspective, having early access to this solution allowed them to bypass struggles with a few x86 specific applications they required that were not available with OpenShift on Power.

■ The process was fairly straightforward and easy to understand. The manual surely helped in giving us some ideas to handle deployments on MAC.

Mauro Sanfilippo
Chief Technology Officer
SME.up

Outcome

The integration of multi-architecture support underscores the Power platform's commitment to adaptability. It empowers users to harness the strengths of different architectures within a unified Red Hat OpenShift Cluster environment. This exciting feature unlocks new possibilities, providing versatility and optimization for composite solutions that span multiple architectures while helping reduce the cost and complexity.



This groundbreaking feature empowers clients to harness their current Red Hat OpenShift cluster to facilitate the deployment of workloads on Power architecture with high availability and co-location advantages. It is financially strategic to run container workloads on Power. Having x86 and Power Worker Nodes in the same cluster simplifies migrating existing x86 applications to Power. While the cost benefits of Power are prominent, MAC allows applications only available on x86 to remain in place, catering to a client's specified needs.

■ This was an excellent experience for us, and I want to thank everyone!

Mauro Sanfilippo
Chief Technology Officer
SME.up

Dive deeper

[Learn how](#) to deploy an open-source e-commerce solution using x86 and Power Worker Nodes with Red Hat OpenShift Multi-Architecture Compute.

Continue learning with the [Getting Started with MAC on Power Guide](#) →



© Copyright IBM Corporation 2024. IBM, the IBM logo, and IBM Power are trademarks or registered trademarks of IBM Corp., in the U.S. and/or other countries.

Red Hat and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Barcelona TechXchange, Jan 2024

smeup LAB

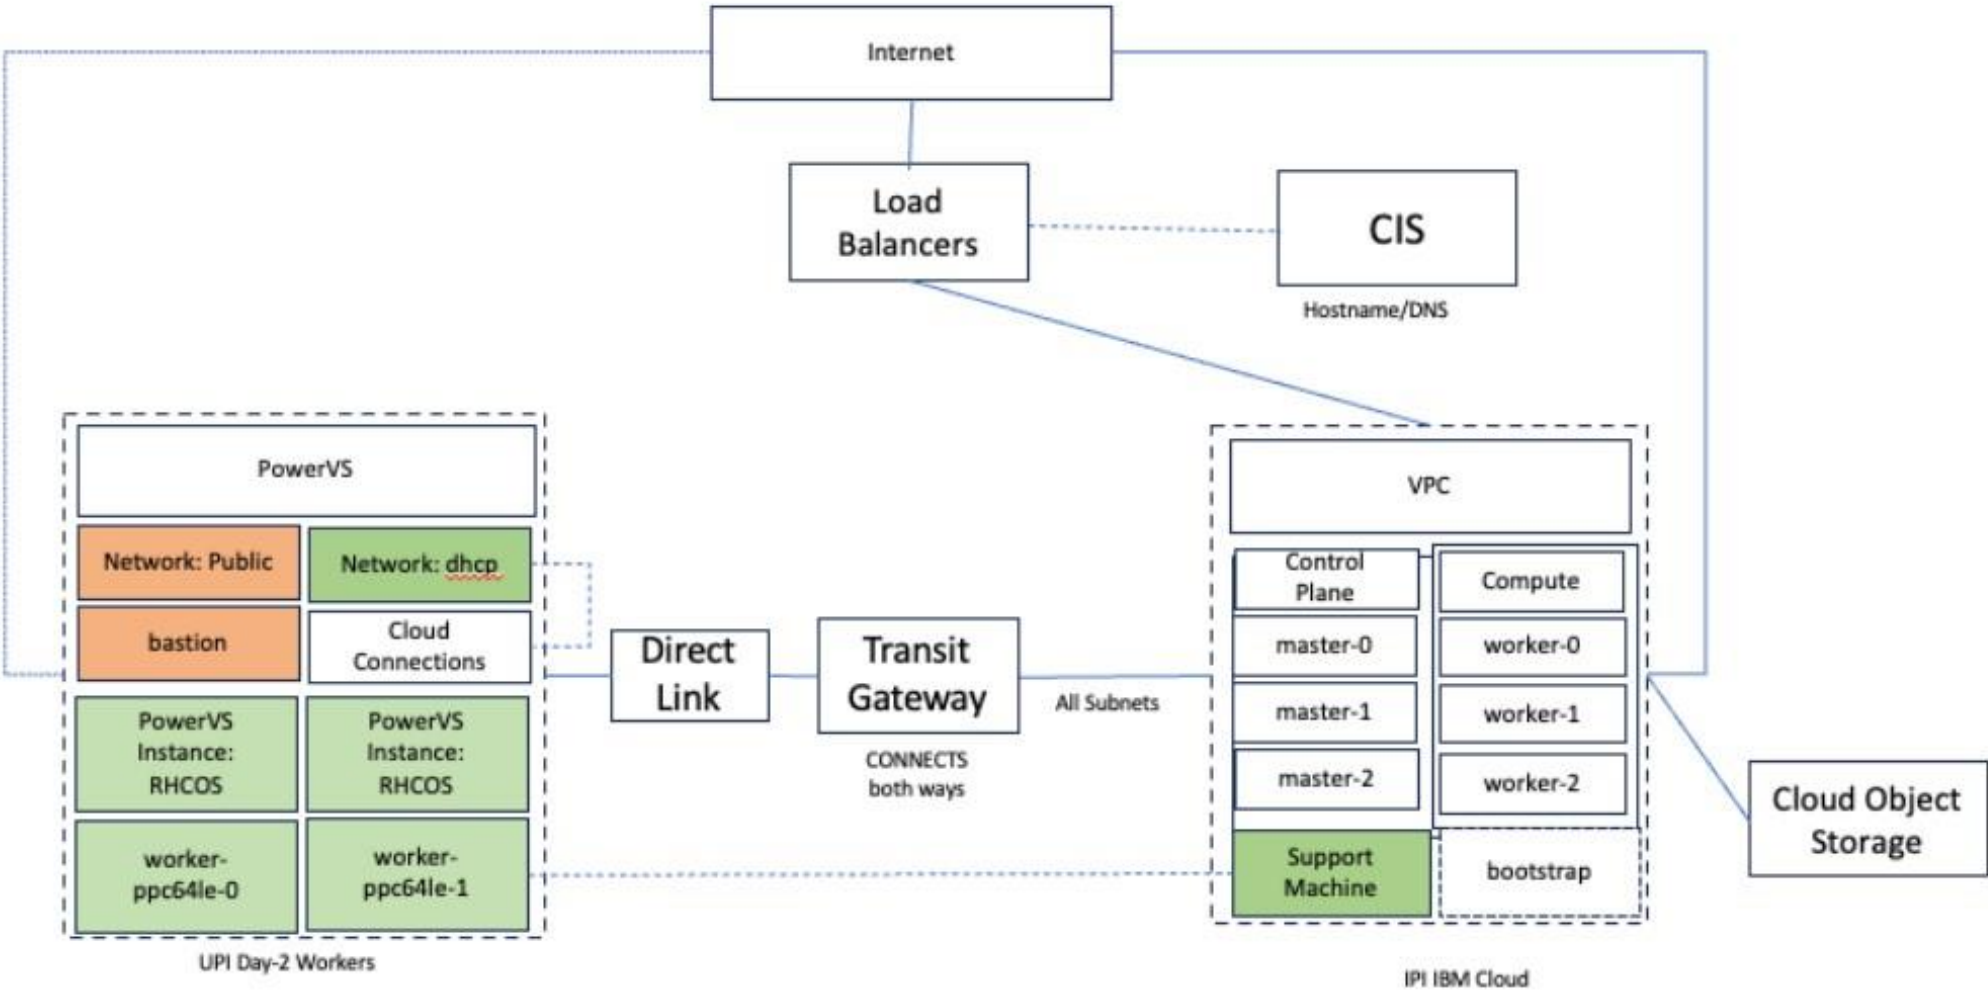
Open Shift Multi-Architecture

Running the same OS Cluster on Intel and Power!

Use cases

- Not supported application runtimes
- Not (yet) supported application versions
- Reuse existing resources

www.smeup.com



Barcelona TechXchange, Jan 2024

smeup LAB

Co-creation Project

data

Paul Chapman · 1^o
Global Power Modernisation Technical Lead

Now you can run x86 & Power Worker Nodes in the same OpenShift Container Platform Cluster with Multi-Arch Compute. ...vedi altro

Vedi traduzione

Agenda

- 01 Why Power?
- 02 What is MAC?
- 03 Why use MAC?
- 04 How to use MAC?
- 05 Early Adoption Program**
- 06 Demo

UKI Brunch & Learn - Red Hat OpenShift - Multi-Architecture Compute

1. Set up MAC Cluster
2. Deploy existing application in MAC
3. Deploy new application
 - Understand **what customer does** in their environment
 - Explore **customer use case/demo**
 - Help customer build **their multi-arch components** (part of MAC Onboarding Essentials Manual in addition to **development support**)
 - Customer **deploys their application** across specific architecture

IBM Team

- Erica Albert
- Paul Bastide
- Paul Chapman
- Geoffrey Pascal

Thanks!

www.smeup.com

© smeup 2024

Barcelona TechXchange, Jan 2024

smeup LAB

Co-creation Project

data

Paul Chapman · 1°
Global Power Modernisation Technical Lead

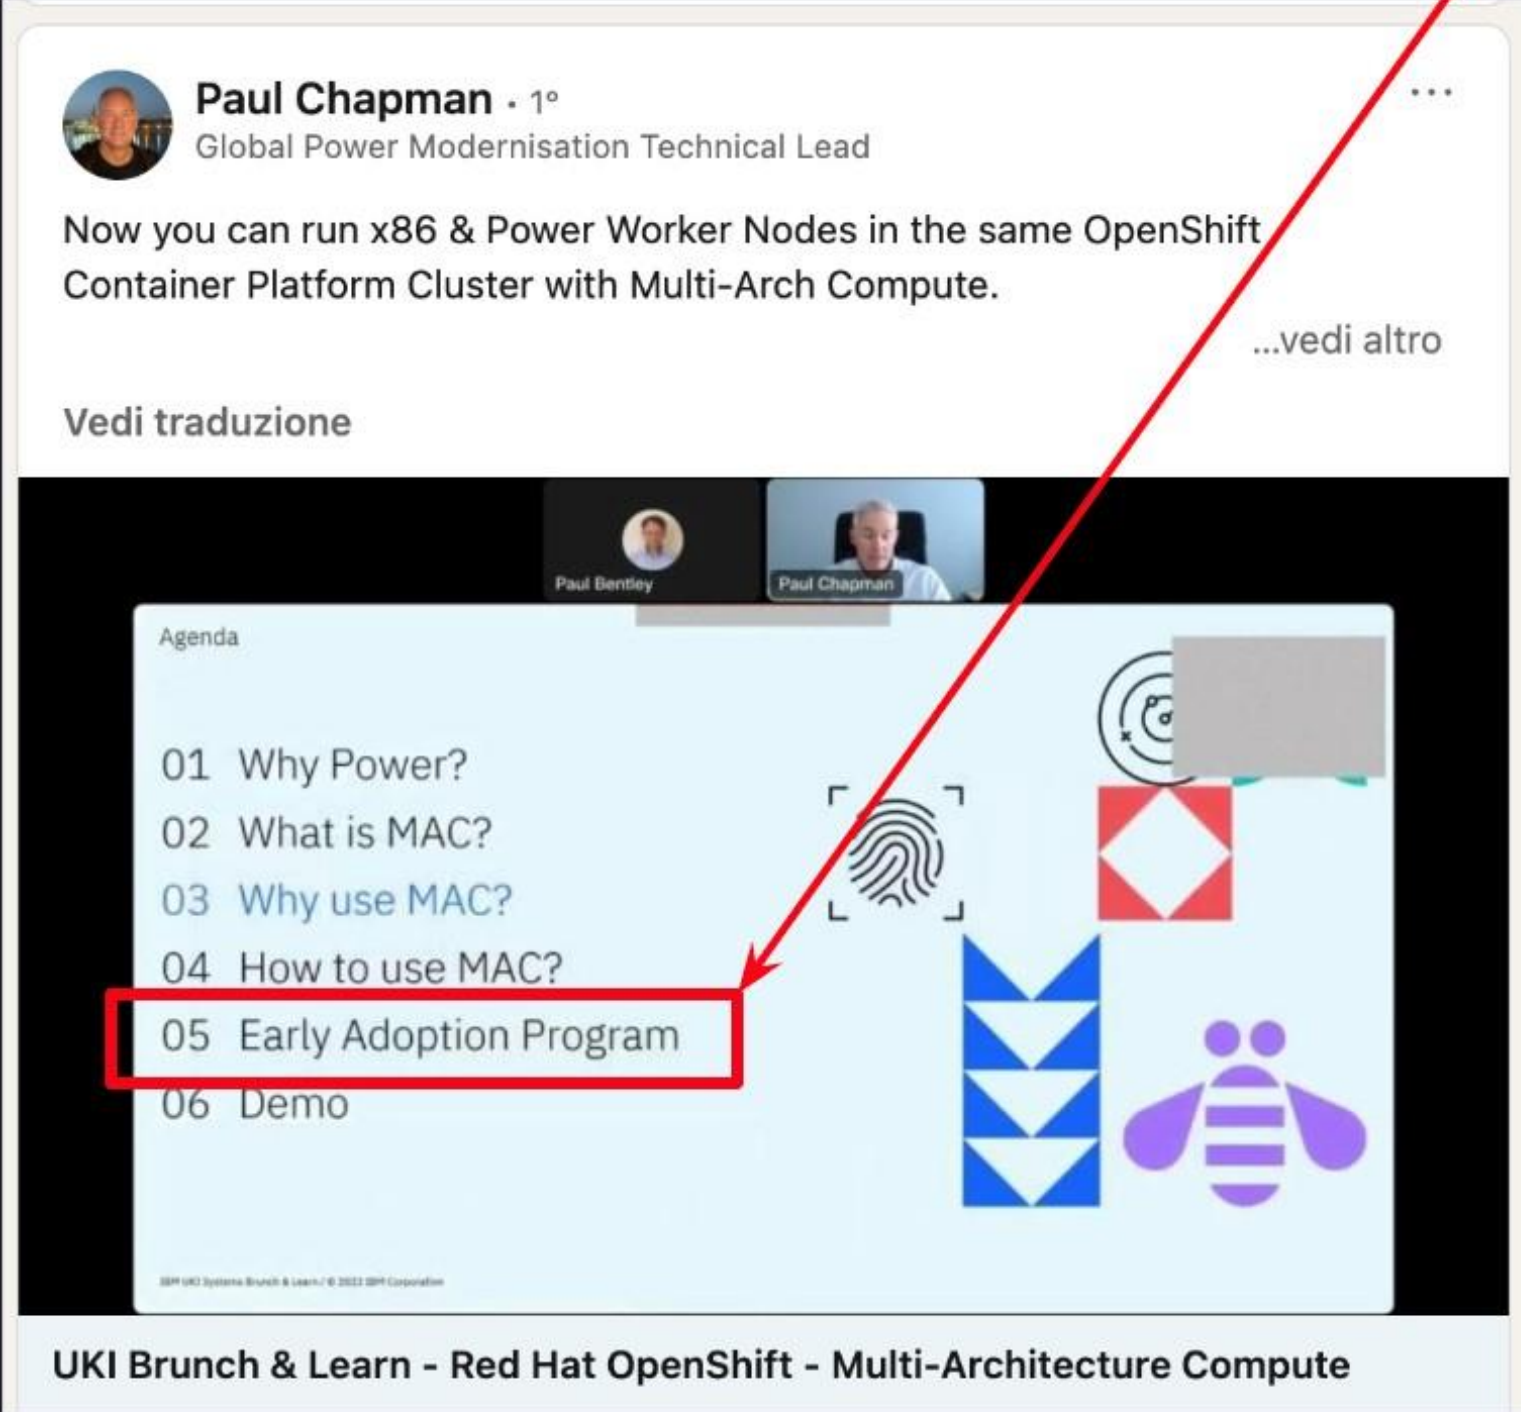
Now you can run x86 & Power Worker Nodes in the same OpenShift Container Platform Cluster with Multi-Arch Compute. ...vedi altro

Vedi traduzione

Agenda

- 01 Why Power?
- 02 What is MAC?
- 03 Why use MAC?
- 04 How to use MAC?
- 05 Early Adoption Program**
- 06 Demo

UKI Brunch & Learn - Red Hat OpenShift - Multi-Architecture Compute



1. Set up MAC Cluster
2. Deploy existing application in MAC
3. Deploy new application
 - Understand **what customer does** in their environment
 - Explore **customer use case/demo**
 - Help customer build **their multi-arch components** (part of MAC Onboarding Essentials Manual in addition to **development support**)
 - Customer **deploys their application** across specific architecture

IBM Team

- Erica Albert
- Paul Bastide
- Paul Chapman
- Geoffrey Pascal

Thanks!

www.smeup.com

© smeup 2024

Agenda

- 01 Multi-Arch Compute History
- 02 IBM Power Strategy
- 03 Prepare for Installation
- 04 Install Multi-Arch sock-shop
- 05 Navigate the sock-shop
- 06 Review Nodes & Pods
- 07 Cordon, Migration Between Architectures
- 08 Multi-Arch Compute Installation
- 09 Public Reference
- 10 [Additional Resources](#)

Multi-Arch Compute - Brunch & Learn

Replay of Paul Chapman introducing and demonstrating Red Hat OpenShift Multi-Arch with Power to IBM UKI Systems, Brunch & Learn attendees.

Now you can run x86, ARM, Power & Z Worker Nodes in the same OpenShift Container Platform Cluster.

Learn how, why and when to add Power to your x86 OpenShift cluster. Watch the recorded live demonstration, plus links to other recordings.

Contact me for help using Power with MAC.

[00:00](#) Introduction

[01:12](#) Request for customer feedback

[02:38](#) Demo - Initial microservice solution build

[05:17](#) Agenda

[06:13](#) Why Power for OpenShift?



[09:25](#) What is Multi-Arch? [12:28](#) Why use Multi-Arch?

[16:07](#) How to use to use Multi-Arch?

[17:07](#) Early Adoption Program - Co-Creation Process

[19:04](#) Demo

[22:37](#) Thank you, Survey & Contact



Demonstration – Deploy Sock Shop to Multi-Arch OpenShift

Demonstration of the deployment of the open-source Sock Shop e-commerce solution using a mix of x86 and Power Worker Nodes with Red Hat OpenShift Multi-Arch Compute

[00:00](#) Intro

[00:14](#) Login to OCP GUI

[00:29](#) OC CLI

[00:37](#) Initiate build of the Sock Shop project to OCP

[01:10](#) Review OCP compute nodes using GUI

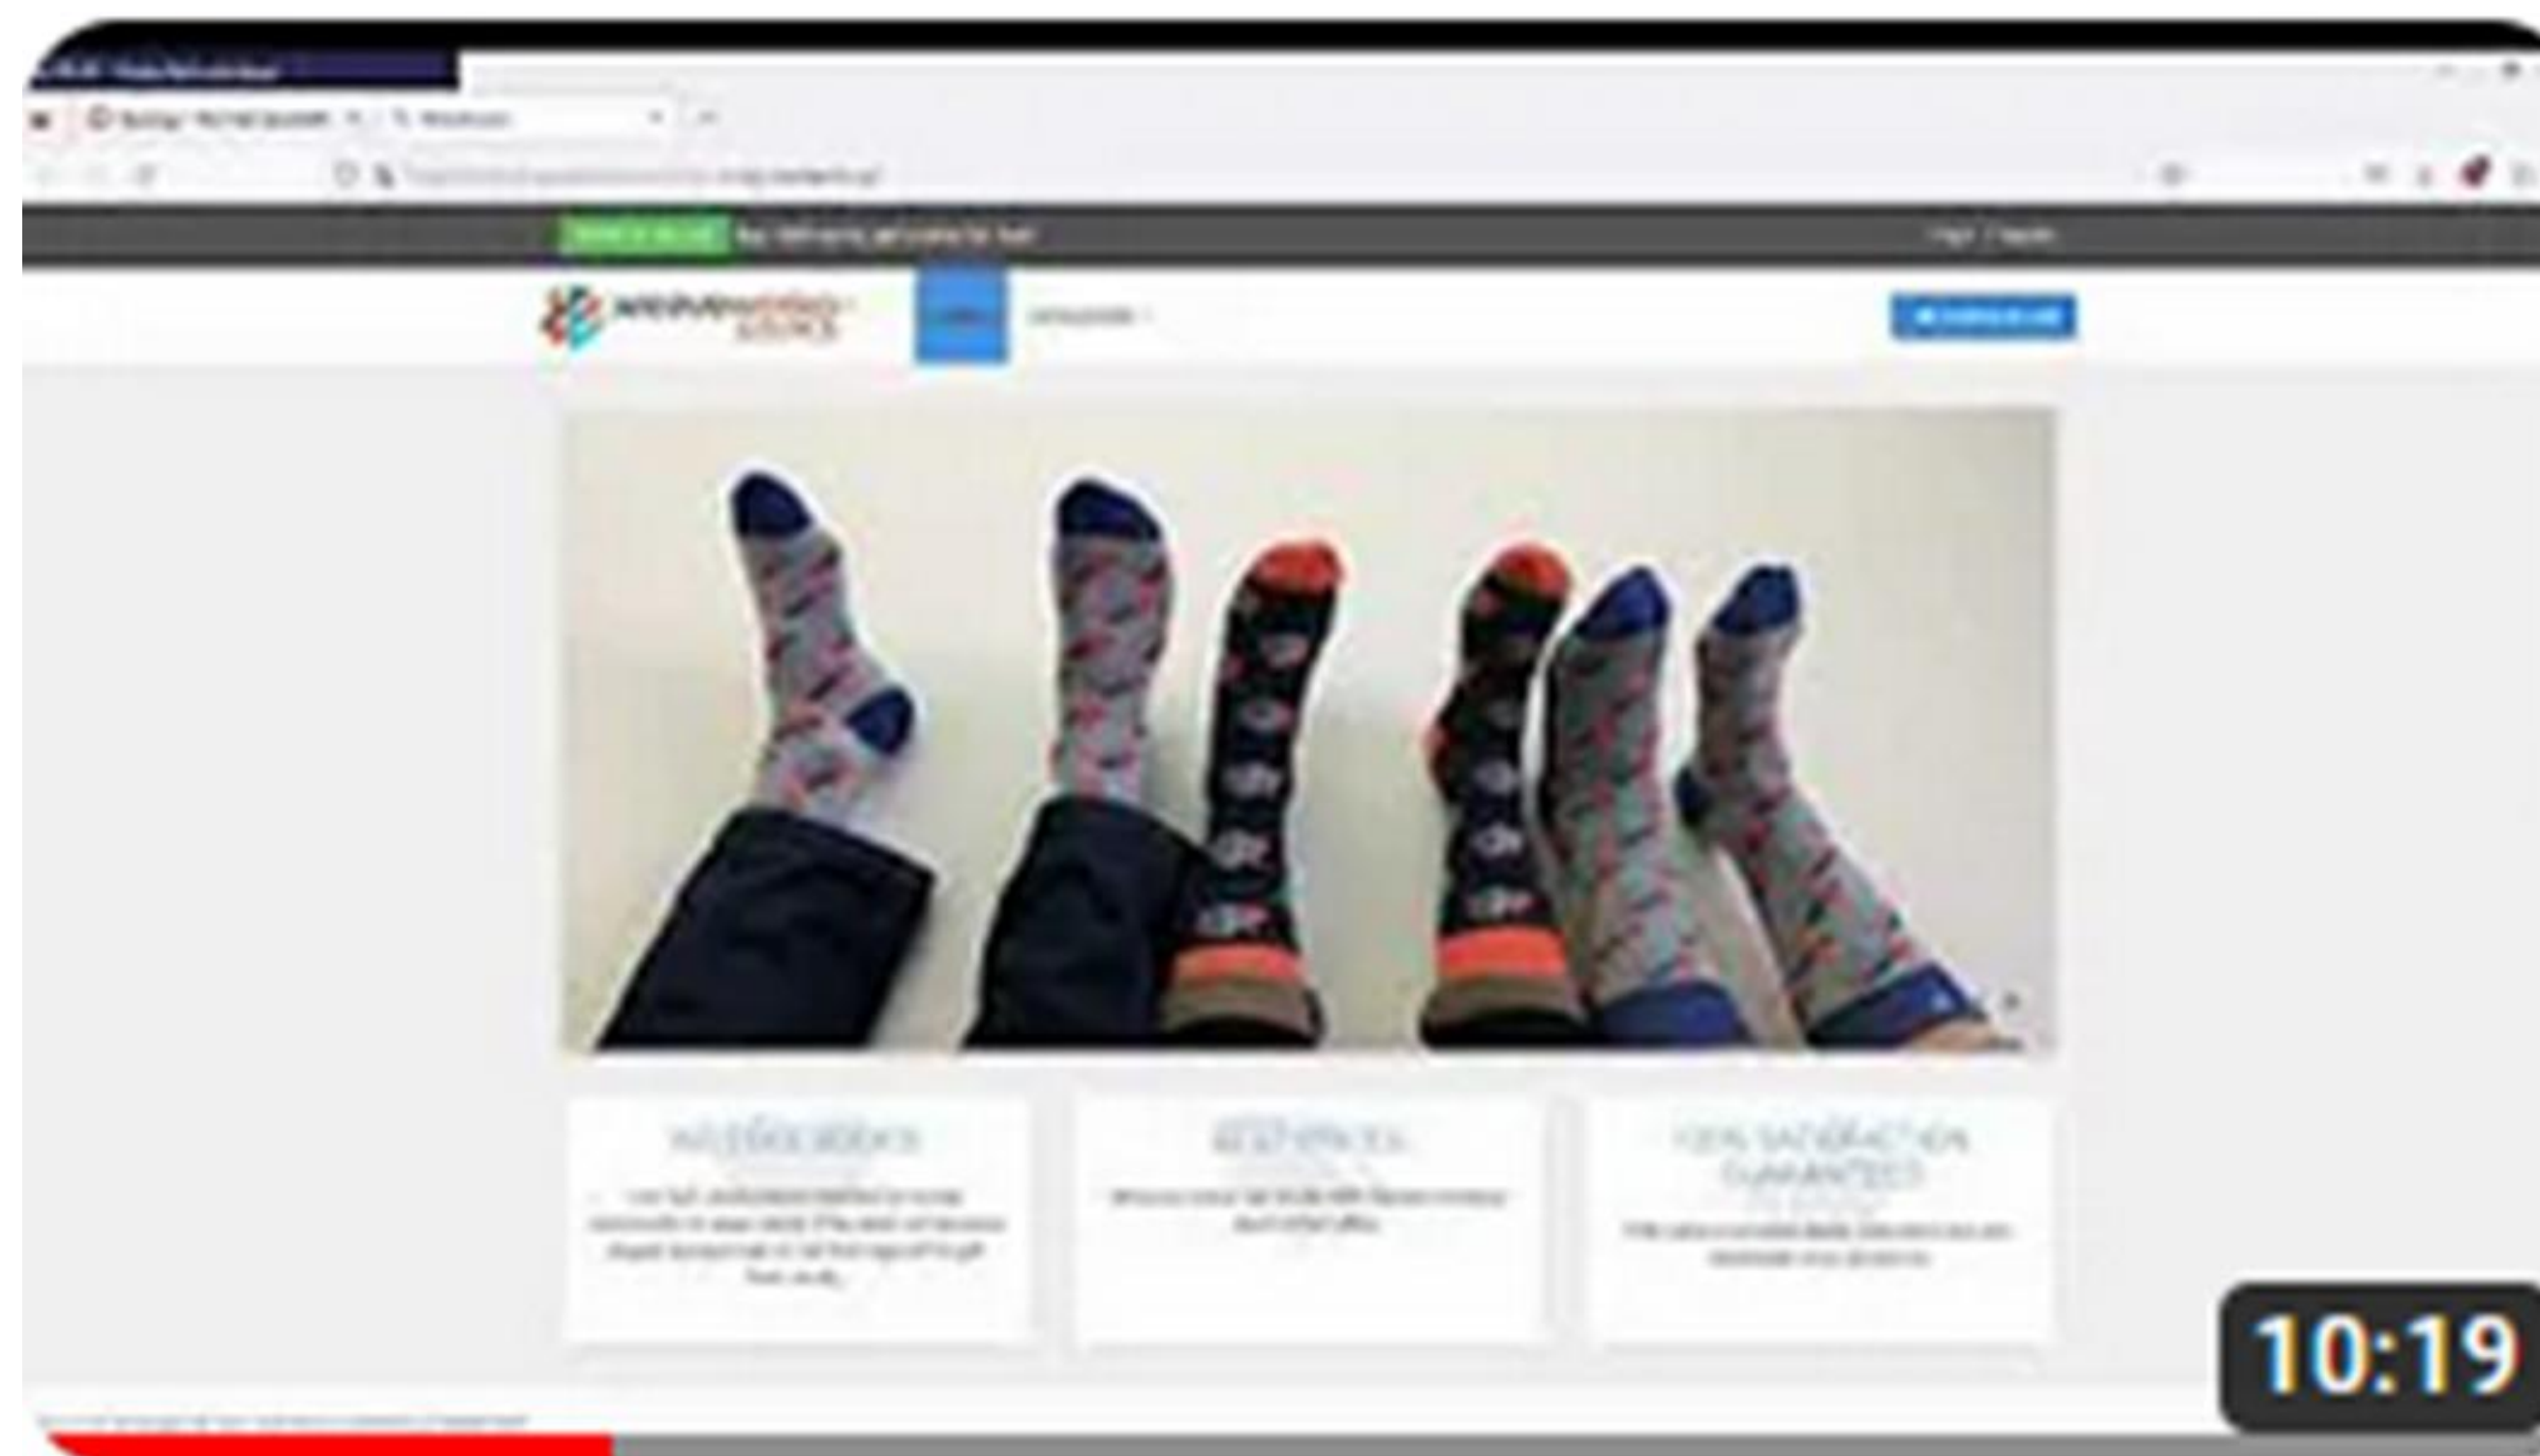
[01:46](#) Review sock-shop project using GUI

[02:32](#) Review micro-service Pod logs using GUI

[03:10](#) Review compute nodes and micro-service Pods using CLI

[05:30](#) Visit the sock shop, register user, review the catalogue, add shipping address and payment details, place and view order using the GUI

[08:41](#) Remove sock-shop project using CLI



OpenShift Multi-Arch Sock Shop Demo

– <https://youtu.be/cas4peunaYQ>

Demonstration – Migrating from x86 to Power Worker Node

This demo features the use of OpenShift Multi-Arch Compute to transfer workloads from an x86 Worker Node to an IBM Power Worker Node.

To begin the migration process, I cordon both x86 and one of the Power Worker Nodes, thereby preventing Pods from being scheduled to them. Subsequently, I delete the Pods on one of the x86 Worker Nodes, which are restarted on the Power Worker Node.

[00:00](#) Sock Shop Installation on Multi-Arch OpenShift

[00:11](#) Sock Shop Web Site

[00:17](#) Sock Shop Nodes

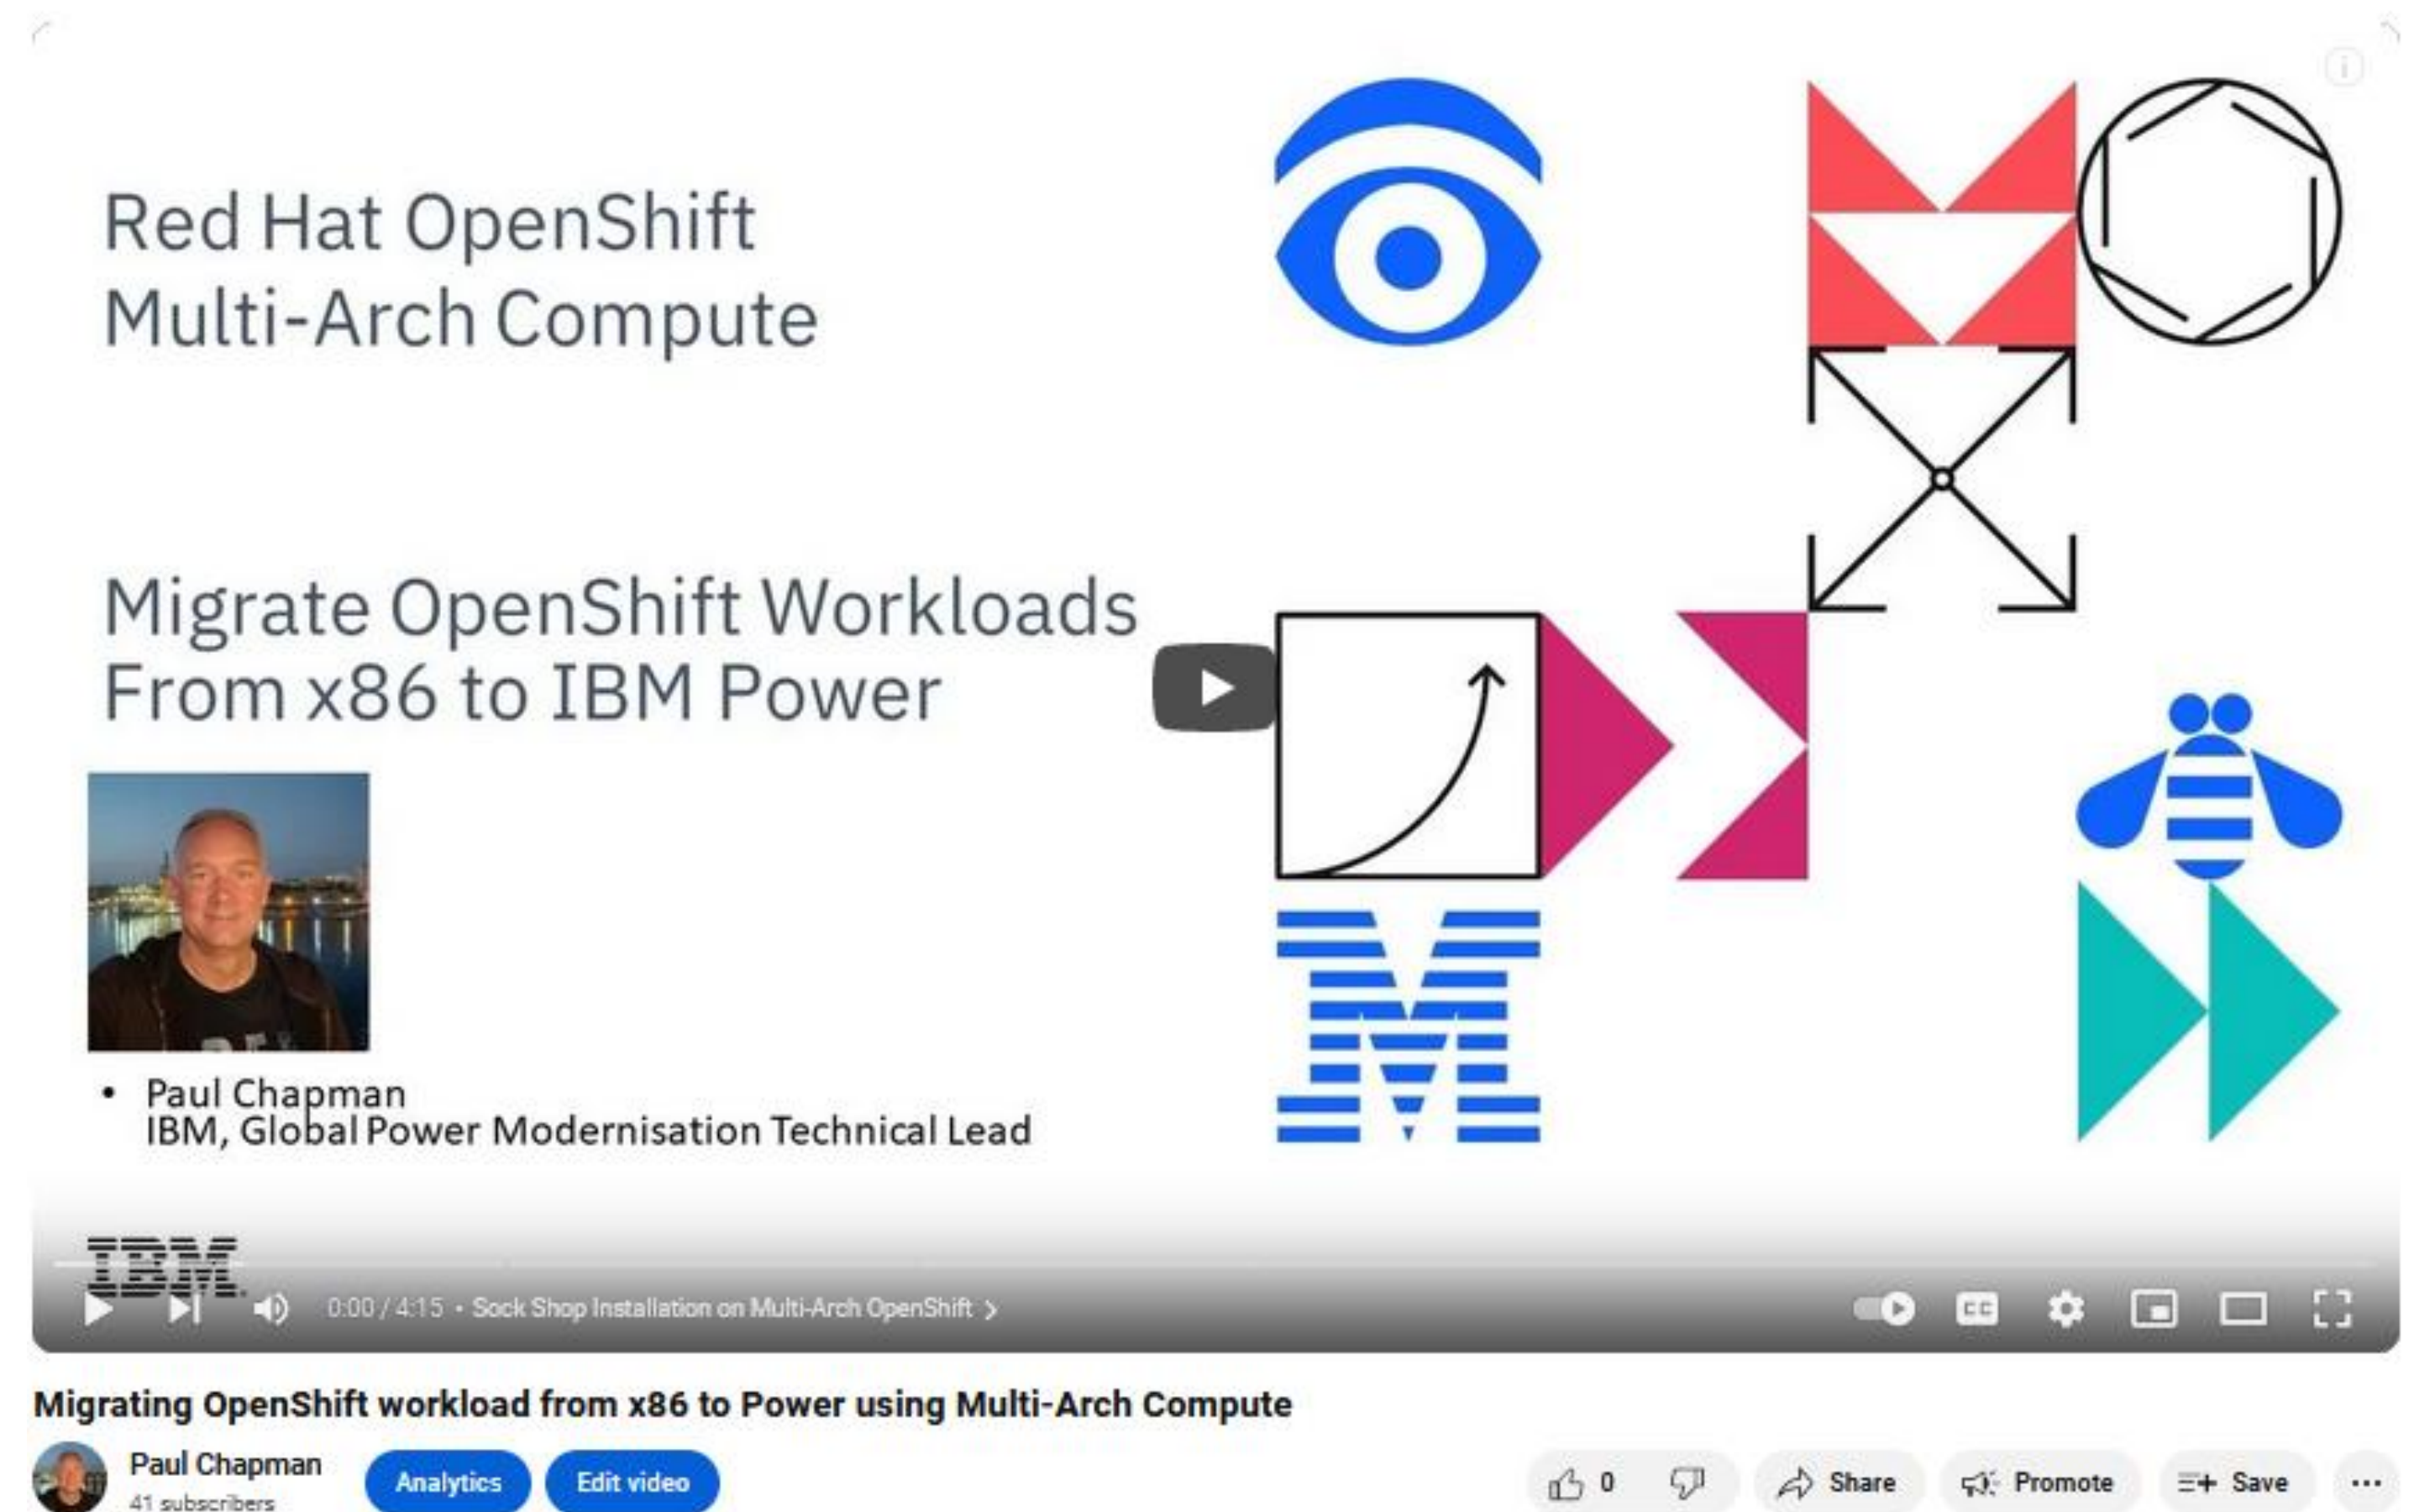
[00:49](#) Sock Shop Pods & Containers

[01:34](#) Cordon Nodes

[01:54](#) Delete x86 Pods

[02:12](#) Pods starting on Power

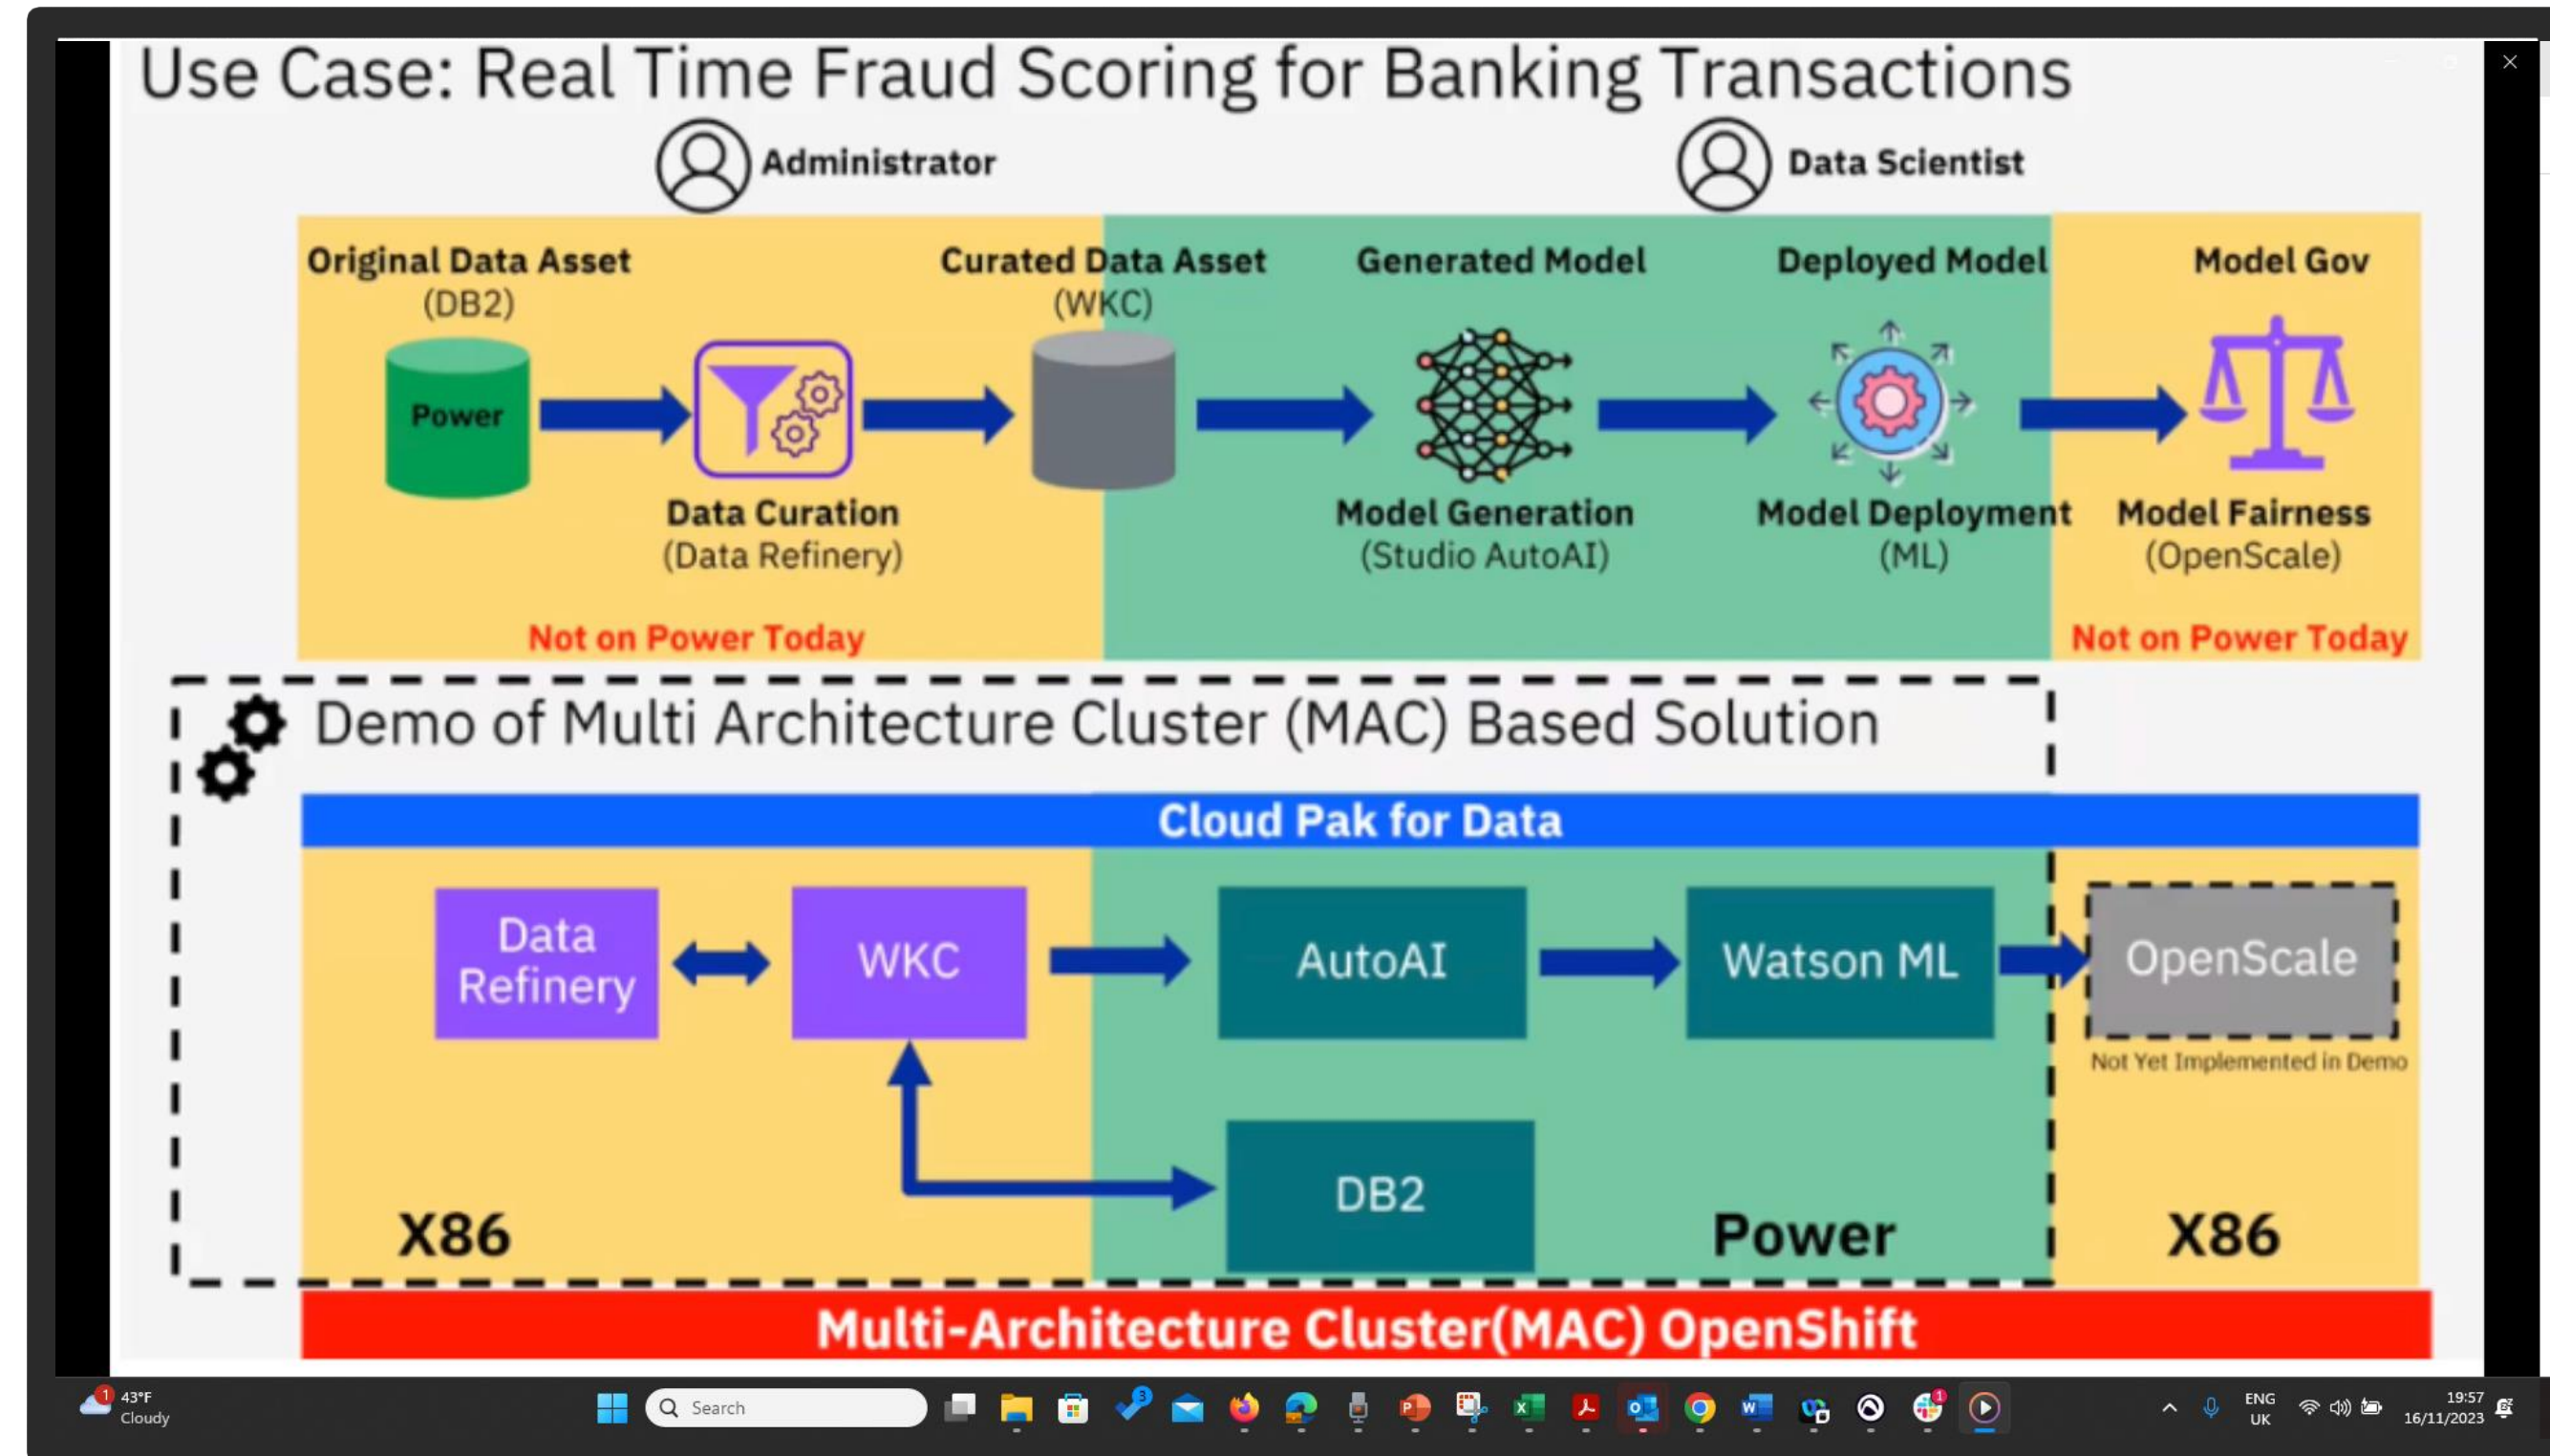
[03:13](#) The Sock Shop is fully functional again



– https://youtu.be/j3SugLx_uQM?si=KBTSrtHxXUhaIaaG



Demonstration - CP4D


Demonstration of real time fraud scoring for banking transactions seamlessly using both x86 and Power Worker Nodes within a single OpenShift Container Platform cluster.

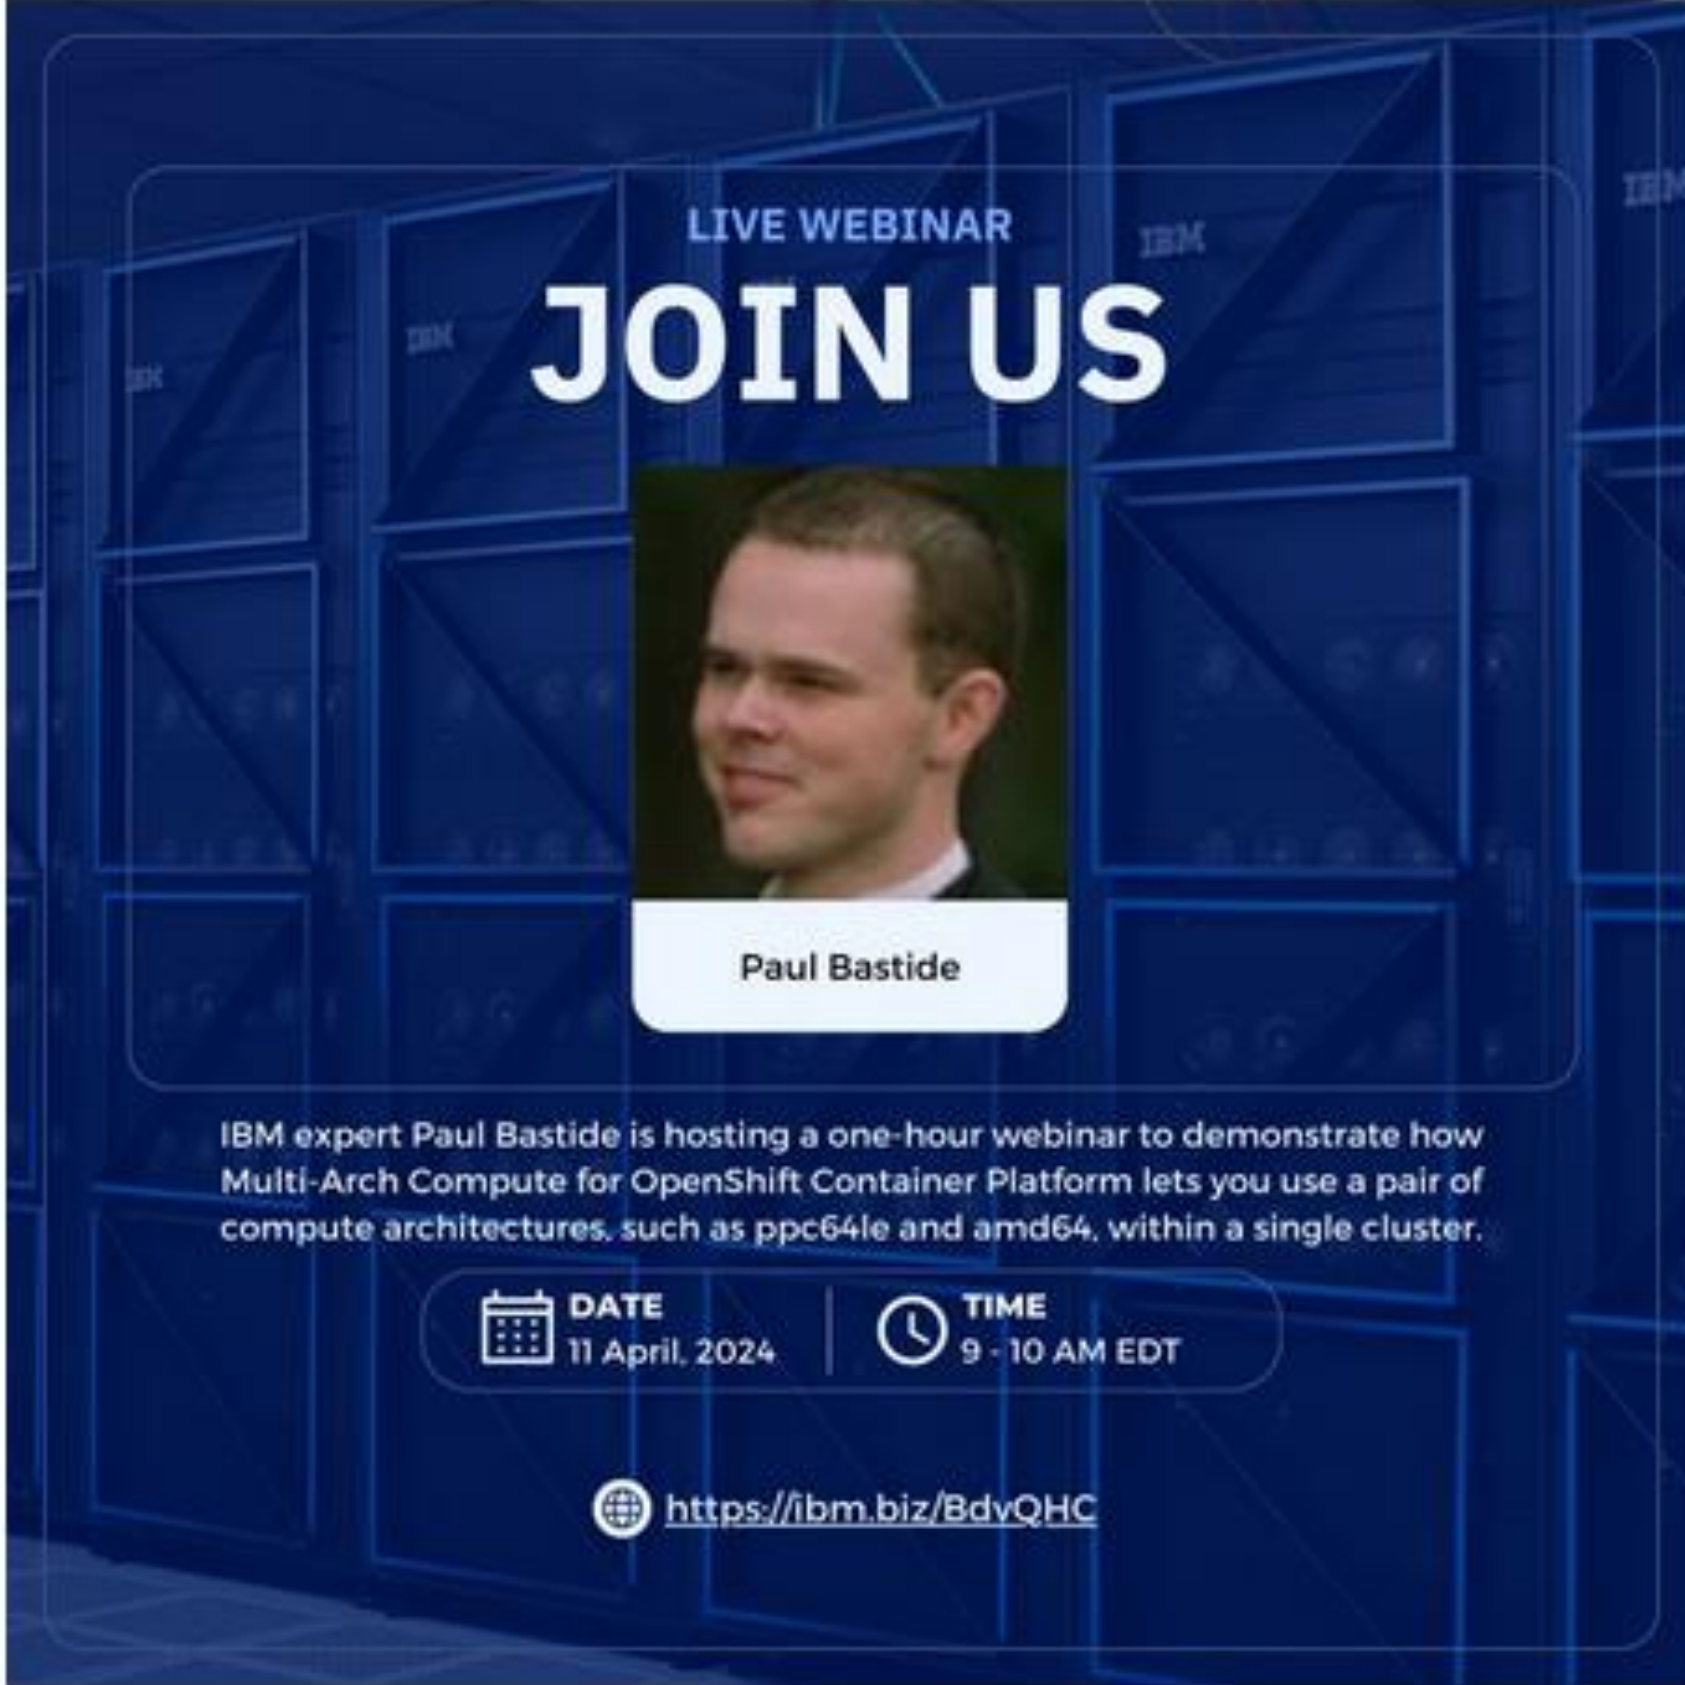


- 5-minute MAC use case demonstration
https://youtu.be/Luo_vbbCt20?si=B5dF1ipiCOFS8sHY


Installation

 **Paul Chapman** • You
Global Power Modernisation Technical Lead
4d • 

Join [Paul Bastide](#), IBM Senior Software Engineer, as he introduces the background behind Multi-Arch Compute and then gets you started setting up, configuring, and scheduling workloads on your  ...see more



LIVE WEBINAR
JOIN US


Paul Bastide

IBM expert Paul Bastide is hosting a one-hour webinar to demonstrate how Multi-Arch Compute for OpenShift Container Platform lets you use a pair of compute architectures, such as ppc64le and amd64, within a single cluster.

DATE 11 April, 2024 | **TIME** 9 - 10 AM EDT

<https://ibm.biz/BdvQHC>

Getting started with Multi-Arch Compute workloads with your Red Hat OpenShift cluster



Paul Bastide
Senior Software Engineer
pbastide@us.ibm.com
LinkedIn: [@paulbastide](#)



Additional Resources

1. [Getting Started ->
https://community.ibm.com/community/user/powerdeveloper/blogs/paul-bastide/2024/02/20/multi-arch-compute-getting-started](https://community.ibm.com/community/user/powerdeveloper/blogs/paul-bastide/2024/02/20/multi-arch-compute-getting-started)
2. [OpenShift Container Platform on IBM Power Systems: Getting Started with Multi-Architecture Compute](#)
3. [Multi-Architecture Compute: Supporting Architecture Specific Operating System and Kernel Parameters](#)
4. [Multi-Architecture Compute: Managing User Provisioned Infrastructure Load Balancers with Post-Installation workers Controlling Pod placement based on weighted node-affinity with your Multi-Arch Compute cluster](#)
5. [Invitation to first MAC Early Adoption and co-creation Program \(90 seconds\)](#)
6. [MAC Demo - Ecommerce OpenSource Microservice Solution \(10 min\)](#)
7. [UKI Brunch & Learn - MAC \(25 min\) - Session Replay: YouTube](#)
8. [OpenShift on IBM Power - Add Intel worker to Power Cluster](#)

1. [Multi-architecture pipelines for IBM Power \(Part 1\)](#)
2. [Multi-arch build pipelines for Power \(Part 2\): Automating multi-arch image builds](#)

Thank You

Paul Chapman
IBM, Global Power Modernization Technical Lead

PaulChapman@uk.ibm.com

Join us again...

More sessions coming...

If you have questions, please contact Paul Bentley bentlep@uk.ibm.com



Register now:



<https://ibm.biz/BdPYQH>


Systems UKI

Watch 100+ recordings:



<https://video.ibm.com/channel/s4Dub4uP9ku>





Notices and disclaimers

© 2024 International Business Machines Corporation.
All rights reserved.

This document is distributed “as is” without any warranty, either express or implied. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.

Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM.

Not all offerings are available in every country in which IBM operates.

Any statements regarding IBM’s future direction, intent or product plans are subject to change or withdrawal without notice.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at: www.ibm.com/legal/copytrade.shtml.

Certain comments made in this presentation may be characterized as forward looking under the Private Securities Litigation Reform Act of 1995.

Forward-looking statements are based on the company’s current assumptions regarding future business and financial performance. Those statements by their nature address matters that are uncertain to different degrees and involve a number of factors that could cause actual results to differ materially. Additional information concerning these factors is contained in the Company’s filings with the SEC.

Copies are available from the SEC, from the IBM website, or from IBM Investor Relations.

Any forward-looking statement made during this presentation speaks only as of the date on which it is made. The company assumes no obligation to update or revise any forward-looking statements except as required by law; these charts and the associated remarks and comments are integrally related and are intended to be presented and understood together.