

# AIOps with OpenShift on IBM zSystems

**Matt Mondics**

Client Technical Enablement Specialist

Hybrid Cloud on IBM zSystems

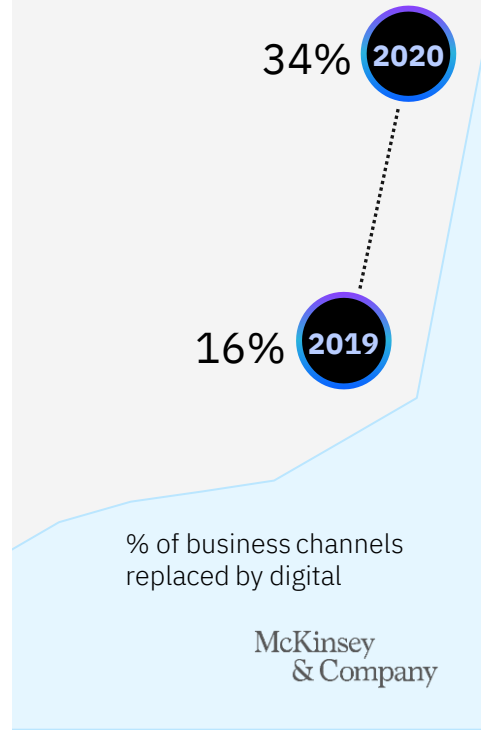
IBM Washington Systems Center

—  
[matt.mondics@ibm.com](mailto:matt.mondics@ibm.com)




## The great digital shift

10 years of digitization  
in under **one** year



## Growing complexity

80% of organizations estimate they have up to 1000 applications in their portfolio today.\*



\*IDC, Worldwide Application Services, 2021,  
Doc # US47588016, April 2021

## Attrition and skills gap

4.3 million people quit jobs in August — about 2.9 percent of the workforce, a pandemic-era record

*The Washington Post, 2021*

1 in 3 workers are now considering leaving their job.

*MagnifyMoney Report, 2021*

# What exactly is AIOps?

## Artificial Intelligence for IT Operations

*The application of data, AI, and machine learning to improve and automate IT operations*


Reduce the operational difficulty in adopting the hybrid cloud model.

AIOps helps to:


- **Collect and aggregate** the ever-increasing volumes of data generated by IT infrastructure components, applications, performance-monitoring tools, and service ticketing systems
- Intelligently shift ‘signals’ out of the ‘noise’ to **identify significant events and patterns** related to application performance and availability issues
- **Diagnose root causes and report them** to IT and DevOps for rapid response and remediation —or, in some cases, **automatically resolve** these issues without human intervention

# IBM AIOps Portfolio


Proactive, continuous application performance



Full-stack, enterprise  
observability



Application resource  
management & cost  
optimization



Proactive incident  
resolution, remediation,  
and avoidance


**INSTANA**  
an IBM Company

**turbonomic**  
an IBM Company


  
IBM Cloud Pak for  
Watson AIOps

# IBM AIOps Portfolio

Proactive, continuous application performance



Full-stack, enterprise  
observability



Application resource  
management & cost  
optimization



Proactive incident  
resolution, remediation,  
and avoidance

**INSTANA**

an IBM Company

**turbonomic**  
an IBM Company




**IBM Cloud Pak for  
Watson AIOps**

# IBM Observability by Instana<sup>®</sup>

Modern applications require modern solutions.

Track **every** interdependency  
from code to end user.



Automate full-  
stack visibility.

Collect accurate  
data in **context**.

Take **intelligent**  
**action**.

# Installing Instana Server

Three paradigms:

1. [Self-hosted on Kubernetes](#)
2. [Self-hosted with Docker containers](#)
3. SaaS – hosted in a public cloud

tl;dr – very simple and quick to install

Instana backend server does not currently support s390x architecture.

1

Install with an Operator

2

1. Install with package manager (yum, apt)
2. run instana init
3. Provide variables when prompted

3

N/A – don't need to install anything

# Installing Instana Server


```
[root@LCSINS01(192.168.176.181) ~ [17:46:30] (0)]# instana init
Setup host environment ✓
? [Please choose Instana installation type] single
? [What is your tenant name?] wsc
? [What is your unit name?] instana
? [Insert your agent key (optional).] 0miTt3d
? [Insert your download key or official agent key (optional).]
? [Insert your sales key] aga1n_0miTt3d
? [Insert the FQDN of the host] LCSINS01.DMZ
? [Where should your data be stored?] /mnt/data
? [Where should your trace data be stored?] /mnt/traces
? [Where should your metric data be stored?] /mnt/metrics
? [Where should your logs be stored?] /var/log/instana
? [Path to your signed certificate file?] /root/cert/tls.crt
? [Path to your private key file?] /root/cert/tls.key
```



# Installing Instana Agent on OpenShift on IBM zSystems

Again, three paradigms:

1. Deploy YAML
2. Deploy Helm chart
3. Deploy operator



**Instana Agent Operator**  
2.0.5 provided by Instana


Install

1

OpenShift

Technology YAML

atsocpd2-s390x OCP on Z

 instana-agent.yaml Copy Download

2

OpenShift

Technology Helm chart

atsocpd2-s390x OCP on Z

```
#!/bin/bash

helm install instana-agent \
  --repo https://agents.instana.io/helm \
  --namespace instana-agent \
  --create-namespace \
  --set openshift=true \
  --set agent.key=qUMhYJxjSv6uZh2SyqTEnw \
  --set agent.downloadKey=qUMhYJxjSv6uZh2SyqTEnw \
  --set agent.endpointHost=LCSINS01.DMZ \
  --set agent.endpointPort=1444 \
  --set cluster.name='atsocpd2-s390x' \
  --set zone.name='OCP on Z' \
  instana-agent
```

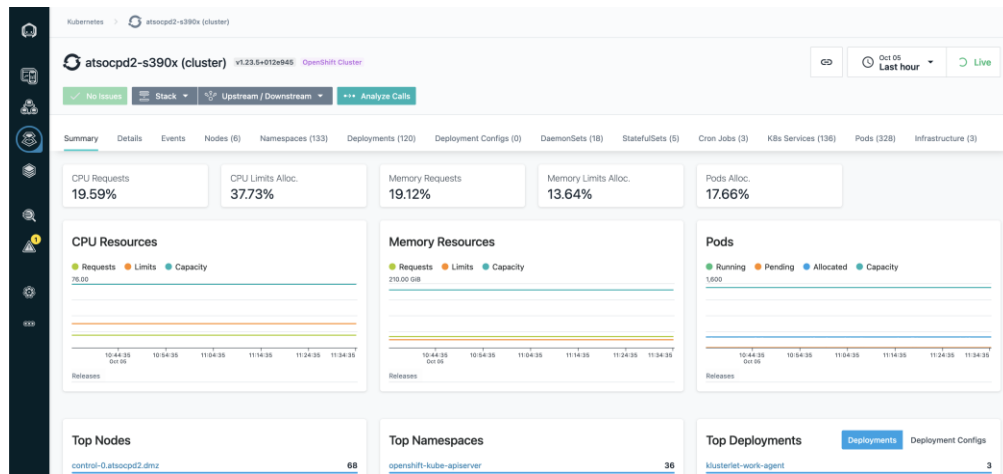
# Installing Instana Agent on OpenShift on IBM zSystems

```
→ atsocpd2-instana-agent$ oc create -f instana-agent.yaml
namespace/instana-agent created
serviceaccount/instana-agent created
secret/instana-agent created
configmap/instana-agent created
daemonset.apps/instana-agent created
clusterrole.rbac.authorization.k8s.io/instana-agent created
clusterrolebinding.rbac.authorization.k8s.io/instana-agent created

→ atsocpd2-instana-agent$ oc get pod -n instana-agent
NAME                                READY   STATUS    RESTARTS   AGE
instana-agent-2plmg                 2/2     Running   0           6m38s
instana-agent-gdp5t                 2/2     Running   0           6m38s
instana-agent-hzd79                 2/2     Running   0           6m38s
```

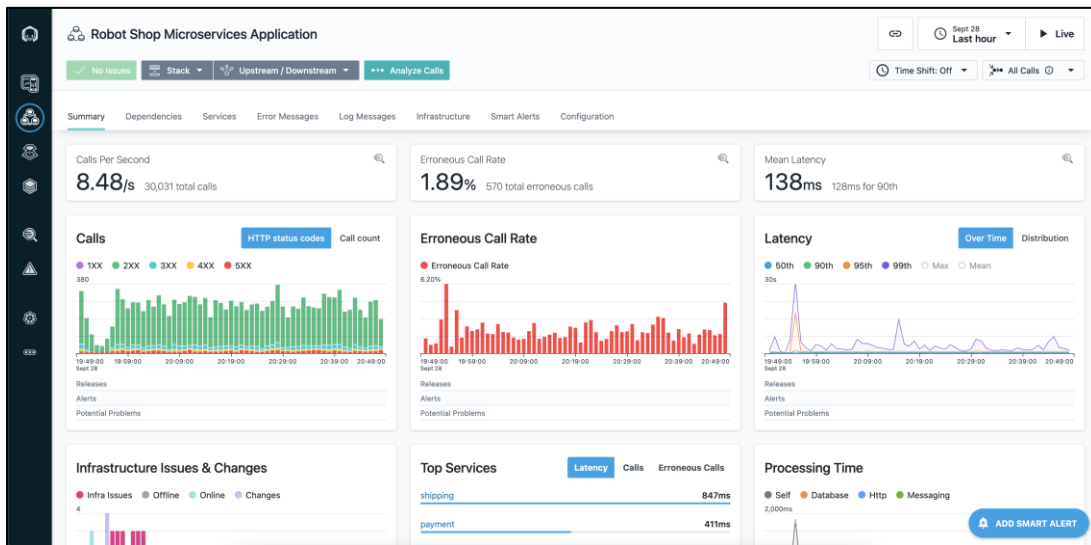
# What can Instana observe with OpenShift on IBM zSystems ?

1. **Cluster health** – overall CPU/memory usage, etcd, scheduler and controller manager health
2. **Cluster events** – all ongoing events in the cluster
3. **Cluster nodes** – CPU/memory usage per node, memory/PID/disk pressures
4. **Kubernetes objects** – namespaces, pods, deployments, services, cron jobs, etc...  
Read more [here](#)
5. **Application traces**, calls, interdependencies, error messages, logs, latency  
Read more [here](#)
6. Built-in **health rules** that can trigger issues and alerts  
Read more [here](#)



# Observing OpenShift on IBM zSystems (and the applications running on it)

1. OpenShift cluster infrastructure
2. Application perspective
3. Events and issues
4. Alerts & Slack

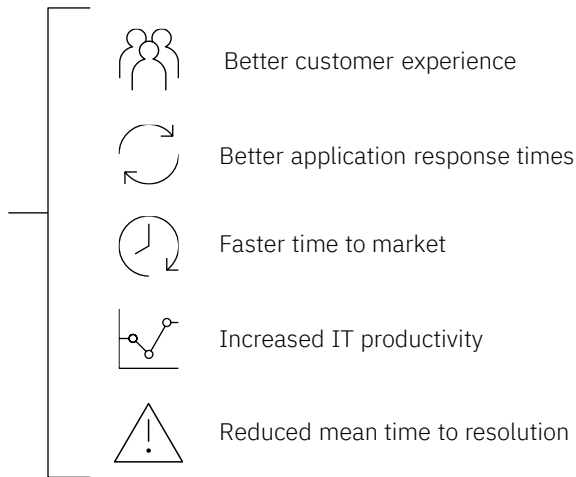
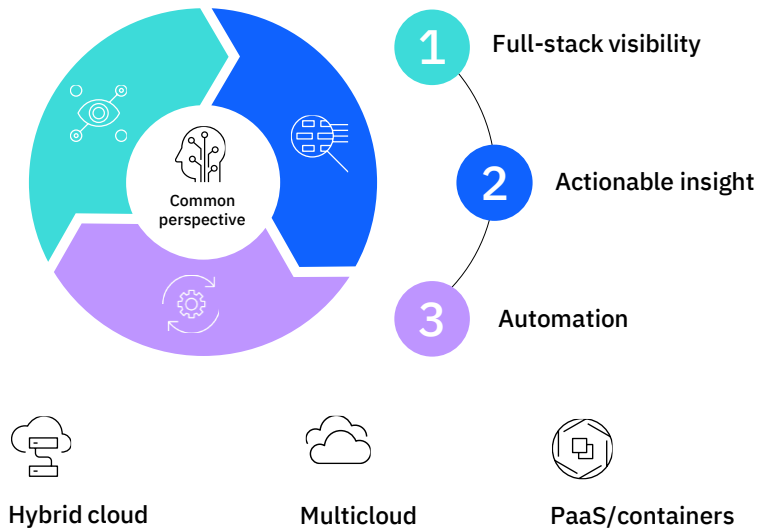


# IBM Turbonomic® Application Resource Management

Continuous application performance while safely reducing cloud costs



471% ROI and  
<6 months payback



# Installing Turbonomic backend

Two paradigms:

1. [Install on OpenShift](#)
2. [Install on a Virtual Machine](#)  
(installs Kubernetes on the VM for you)

Turbonomic *server* does not currently support IBM zSystems architecture

However you can *manage-to* OpenShift on IBM zSystems

1

Install with an Operator

2

1. Install vCenter or Hyper-V
2. Deploy Turbonomic components

# Installing Turbonomic backend



## Turbonomic Platform Operator

42.16.0 provided by Turbonomic, Inc.

✕

Install

**Latest version**

42.16.0

**Capability level**

☒ Basic Install

☐ Seamless Upgrades

☐ Full Lifecycle

☐ Deep Insights

☐ Auto Pilot

### Realtime Decision Automation for Multicloud Applications

Turbonomic Workload Automation for Multicloud simultaneously optimizes performance, compliance, and cost in real-time. Workloads are precisely resourced, automatically, to perform while satisfying business constraints:


- Continuous placement of workload across multiple clouds both on-prem and public clouds providers.
- Continuous scaling for applications and the underlying infrastructure.

It assures application performance by giving workloads the resources they need when they need them.

# Configuring OpenShift on IBM zSystems as a Turbonomic “Target”

Deploy [Kubeturbo operator](#) on the target OpenShift cluster

Can also integrate applications managed by an observability solution i.e. Instana, AppDynamics, Dynatrace, New Relic

 **Kubeturbo Operator**  
8.6.4 provided by Turbonomic, Inc. ✕

Install

**Latest version**  
8.6.4

**Capability level**  
☒ Basic Install  
☐ Seamless Upgrades  
☐ Full Lifecycle  
☐ Deep Insights  
☐ Auto Pilot

**Source**

### Application Resource Management for Kubernetes

Turbonomic AI-powered Application Resource Management simultaneously optimizes performance, compliance, and cost in real time. Software manages the complete application stack, automatically. Applications are continually resourced to perform while satisfying business constraints.

Turbonomic makes workloads smart – enabling them to self-manage and determines the specific actions that will drive continuous health:

- Continuous placement for Pods (rescheduling)
- Continuous scaling for applications and the underlying cluster.

It assures application performance by giving workloads the resources they need when they need them.

Top Business Applications

Global Environment

?

⋮

<



# What actions can Turbonomic recommend and automate?

1. Vertically scale workloads  
(uses more historical data than other autoscalers like OCP VPA)
2. Horizontally scale workloads (to meet SLO)
3. Provision or suspend cluster nodes  
(recommend only on s390x)
4. Move pods to different nodes  
(unique to Turbonomic)

[More details](#)

1

Modify CPU requests and limits, equivalent to:

```
oc set resources deployment foo \  
--limits=cpu=200m,memory=512Mi \  
--requests=cpu=100m,memory=256Mi
```

2

Increase or decrease number of pods, equivalent to:

```
oc scale deployment foo \  
--current-replicas=1 \  
--replicas=2
```

3

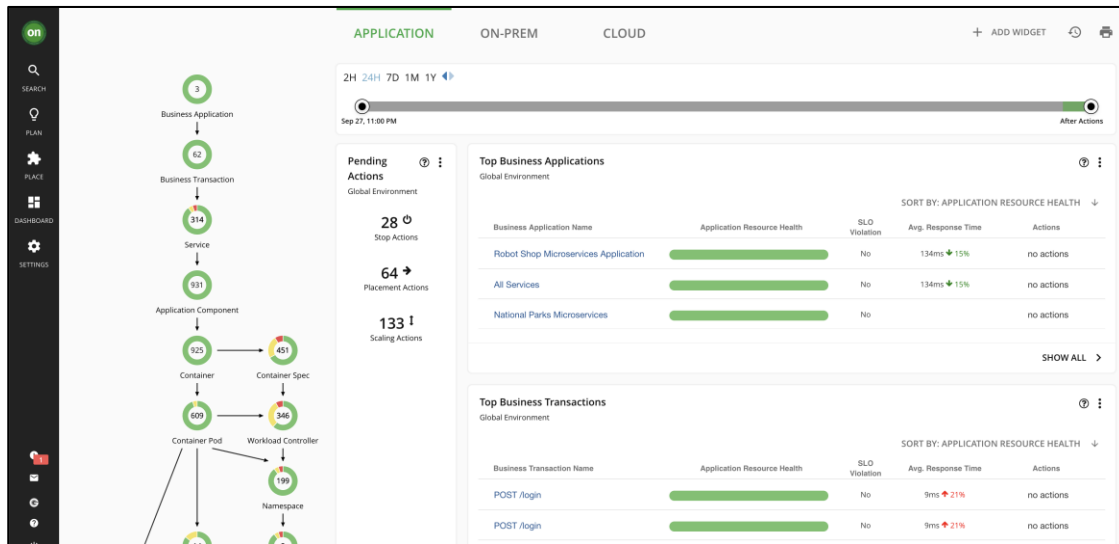

Cannot *automatically* scale nodes on s390x –  
however you can manually scale when *recommended*

4

Spin up new pod on target node -> ensure it is up and running and passes health checks -> delete original pod

# Managing OpenShift on IBM zSystems resources with Turbonomic

1. Topology
2. Pending Actions
3. Optimize container cluster plan
4. Create an automated policy
5. Dashboards



# Caveats, limitations, things to think about

- 1. Managing application resources that are already managed by operators**  
Must create an [Operator Resource Mapping \(ORM\)](#) for each operator  
ORMs must be developed by each solution team (Cloud Paks for example)
- 2. Pods that use PVCs that are RWO (read-write-once) cannot be moved without interruption**  
Delete original pod (frees up PVC) -> spin up new pod on target node -> new pod binds to PVC
- 3. Don't trust Turbonomic to automatically manage your cluster/application resources?**  
Perform the recommended actions manually, add more and more automation over time...  
Similarly, nothing is stopping you from manually taking actions that cannot be done automatically (i.e. suspend or provision cluster nodes)

# More resources

1. [Instana product page](#)
2. [Instana documentation](#)
3. [Turbonomic product page](#)
4. [Turbonomic documentation](#)
5. [CP4WAIOps product page](#)
6. [CP4WAIOps documentation](#)
7. [Video: Monitoring Kubernetes and OpenShift with Instana](#)
8. [Video: Turbonomic 8 Walkthrough](#)
9. [Datasheet: IBM Turbonomic Application Resource Management for OpenShift](#)

# Thank you

Matt Mondics

Client Technical Enablement Specialist

Hybrid Cloud on IBM zSystems

—

matt.mondics@ibm.com



© Copyright IBM Corporation 2021. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., 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 at [Copyright and trademark information](#).

