Avanseus
July 2021

STRICTLY CONFIDENTIAL

avanseus

Cognitive Assistant For Networks Release 5.5 Major Features

Artificial Intelligence for Telecom, Enterprise & IoT Applications- Release 5.5, July 2021

Gartner Cool Vendor 2018

Cool Vendor in Communications Service Provider Business Operations

The Gartner, "Cool Vendors in Communications Service Provider Business Operations", was written by analysts Norbert Scholz | Sylvain Fabre | Jouni Forsman | Ian Keene | Martina Kurth | Wm. L. Hahn | Peter Liu | Ramesh Marimuthu | Amresh Nandan | Michael Porowski | Kosei Takiishi and published on 24 September 2018. The Gartner Cool Vendor Logo is a trademark and service mark of Gartner, Inc., and/or its affiliates, and is used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in our research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

CAN 5.5 Major Features

1 Service now integration

2 Kafka integration

3 Kubernetes based deployment

4 Advanced RoE & Ticket matching

5 Tomcat clustering for session management in CAN & CAS

6 REST API for prediction delivery

7 Elasticstack integration in Openshift

8 Parser Screen For Work Order Processing

9 Horizantal pod autoscaling – HPA

10 Web Security porting from Apache to Ngnix

11 UI Porting with React Framework

12 Memcached tool integration for centralized cache

13 Integration with Istio service mesh

14 Prediction as a Service

15 Helm Charts for easy installation

16 Integration with Monitoring softwares (Kiali, Promotheus and Grafana)

17 Re-scheduling of trigger based on UI cron pattern configurations

CAN 5.5 -ServiceNow Integration

- ServiceNow is a ticketing/incident management software for booking trouble tickets and managing them.
- CAN 5.5 integrates ServiceNow over the integration gateway interface
- Enables new predictive ticket creation and assignment
- Enables updating of existing predictive tickets
- Enables Real-time updates from CAN 5.5 to ServiceNow and vice versa
- Maps all essential field automatically & also an option to customize, extra fields depending on user requirement are available

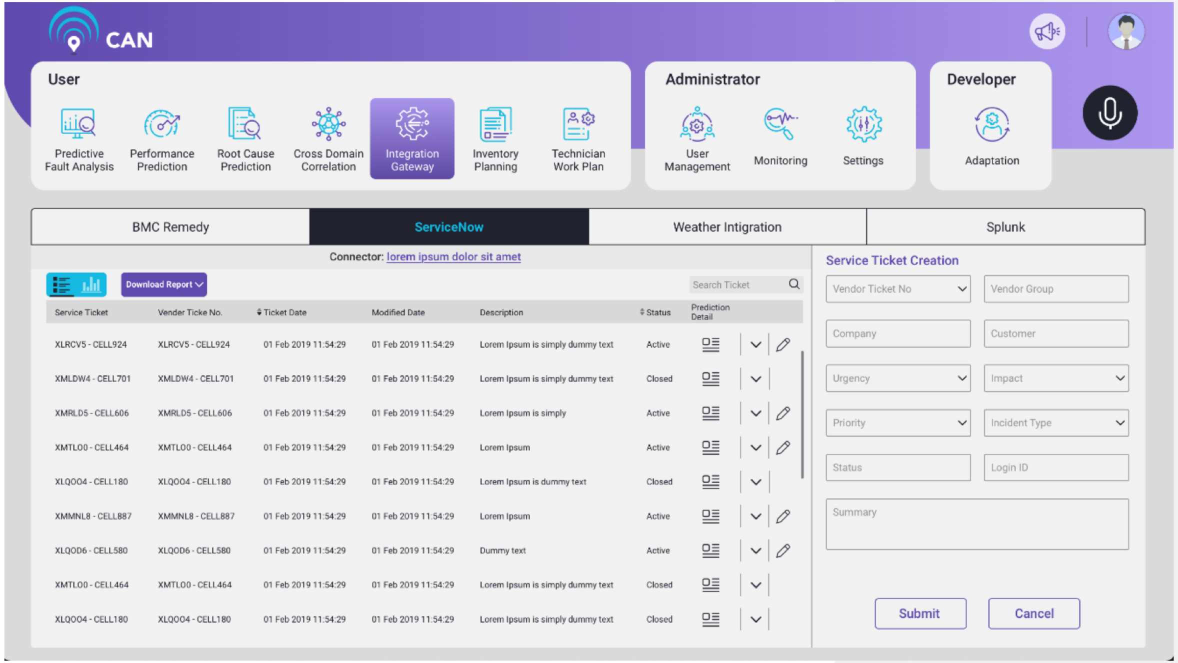

Connector: lorem ipsum dolor sit amet

Service Ticket	Vender Ticke No.	Ticket Date	Modified Date	Description	Status	Prediction Detail
XLRCV5 - CELL924	XLRCV5 - CELL924	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is simply dummy text	Active	
XMLDW4 - CELL701	XMLDW4 - CELL701	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is simply dummy text	Closed	
XMRLD5 - CELL606	XMRLD5 - CELL606	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is simply	Active	
XMTLO0 - CELL464	XMTLO0 - CELL464	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum	Active	
XLQOO4 - CELL180	XLQOO4 - CELL180	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is dummy text	Closed	
XMMNL8 - CELL887	XMMNL8 - CELL887	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum	Active	
XLQOD6 - CELL580	XLQOD6 - CELL580	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Dummy text	Active	
XMTLO0 - CELL464	XMTLO0 - CELL464	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is simply dummy text	Closed	
XLQOO4 - CELL180	XLQOO4 - CELL180	01 Feb 2019 11:54:29	01 Feb 2019 11:54:29	Lorem Ipsum is simply dummy text	Closed	

CAN 5.5 – ServiceNow Integration Architecture

Salient Features

- Seamless Connector Application to interconnect with ServiceNow
- Seamless transition from prediction to ServiceNow tickets.
- Ticket generation can be automated on threshold configuration.
- Direct ticket booking from Prediction list
- Ticket editing option from CAN
- Ticket searching option from CAN
- Changes on open tickets will be visible in CAN in real time
- Ticket report downloading option
- Ticket closure and Engineer assignment happens from ServiceNow
- Option to use customer specific fields for fault management is available

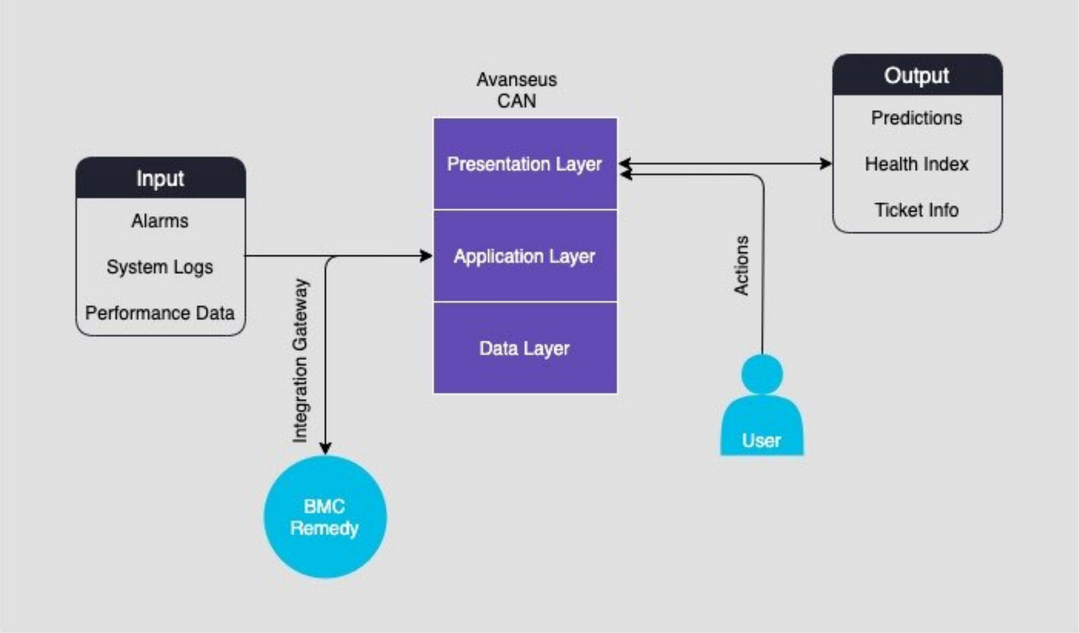

CAN 5.5 - Kafka Integration

- Kafka framework allows processing of streaming data
- It is a high throughput, low latency platform to handle real time data feeds
- Enables decoupling of data producers and consumers
- CAN integration with Kafka broker optimizes customer operations of sending alarms, tickets & performance counter data in a streaming channel.
- Kafka integration enables CAN to consume the incoming data through kafka broker without actually connecting to customer systems (producer) directly.
- Kafka Streaming interface allows CAN to subscribe to it and digest the data in realtime.

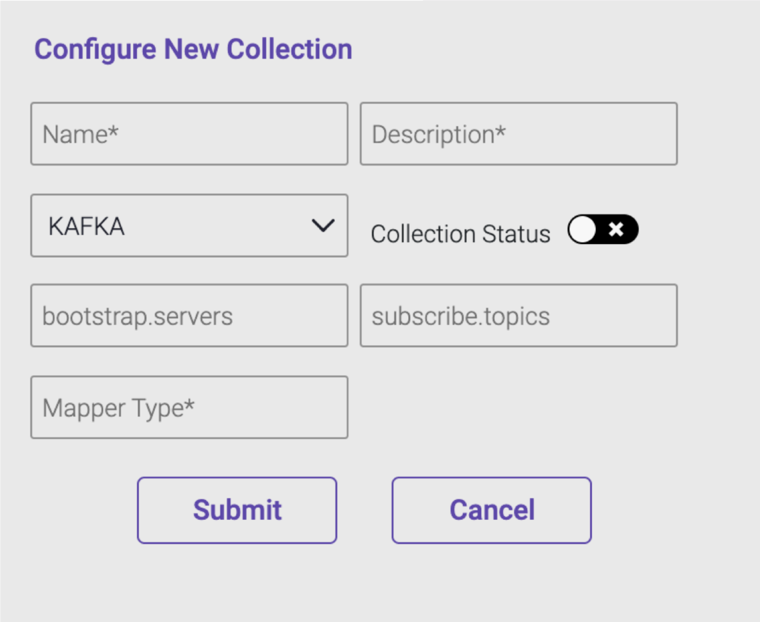

Configure New Collection

Name*	Description*
KAFKA	Collection Status
bootstrap.servers	subscribe.topics
Mapper Type*	

CAN 5.5 – Kafka Integration Architecture

Salient Features :

- Seamless Streaming Interface to interconnect with Kafka broker
- Option to subscribe the required data stream for automated and realtime data access
- Optimized operations for customer as no changes need to be carried out at customer end for any addition of data consumers.
- Improved efficiency through load balancing and partitioned storage
- Filtering of specific data or specific topics for streaming can be enabled at kafka cluster based on CAN requirement enabling optimization of operations

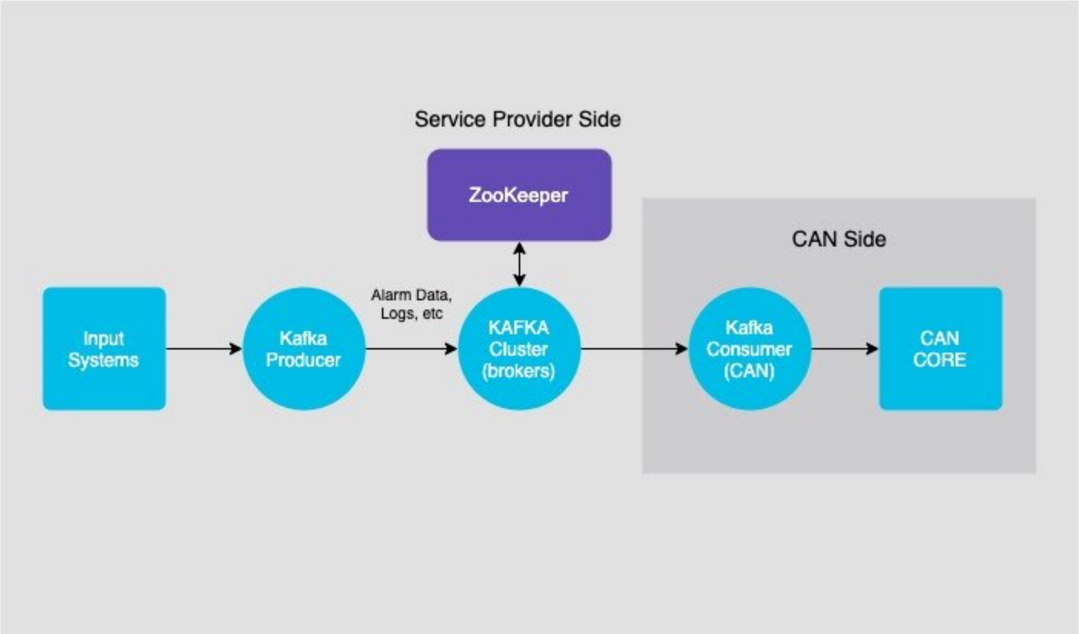

CAN 5.5 - Kubernetes based deployment

- Enables CAN compatible to Cloud Native Environments
- CAN 5.5 is compatible with Kubernetes tools to enable automation of deployment, scaling and management of CAN pods
- VM based deployment is replaced by Kubernetes based deployment in CAN 5.5.
- Enables CAN software to use Kubernetes services like load balancing, self healing, auto scaling etc.,

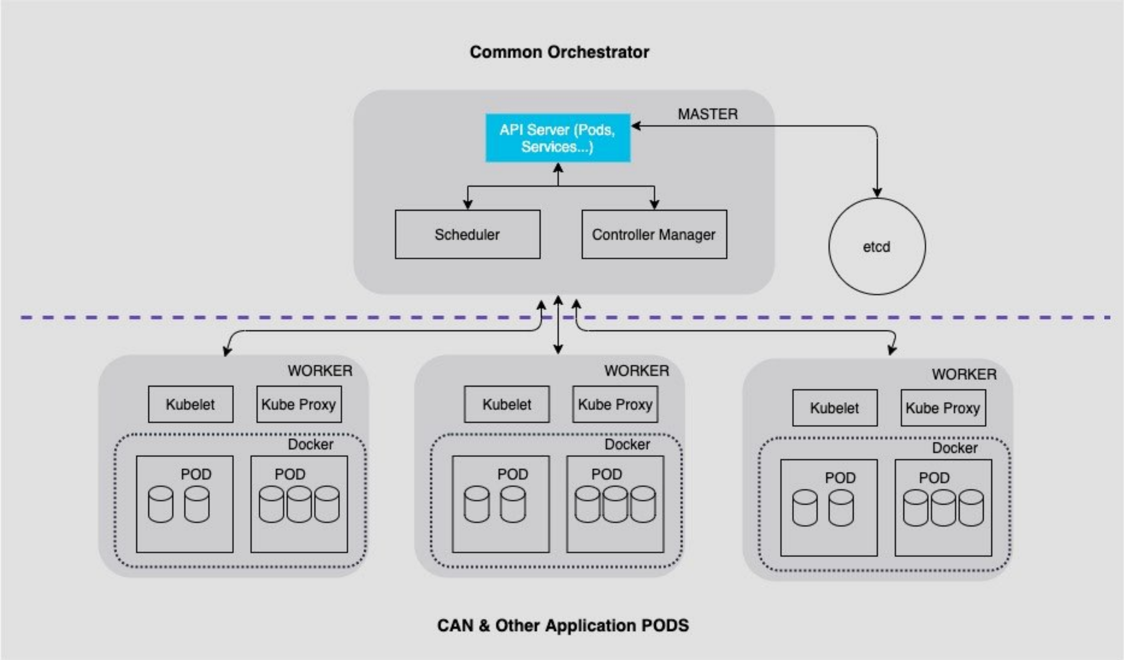

CAN 5.5 - Advanced RoE and Ticket Matching

- Return of Effort (RoE) is an existing feature of CAN where a particular subset of predicted faults that are more impactful or likely to happen are highlighted
- This is an autonomous feature that enables ease of operation & saves precious time of customers by providing "ready for action" reports
- Advanced RoE provides more accurate and meaningful predictions as more parameters like ticket correlation, rarity of alarms, work orders issued against faults etc are taken into consideration in addition to older policies
- Advanced RoE also have control over the number of predictions that can be selected using the RoE configurations
- Users are provided with UI support to enable and disable RoE policies and also to add and delete policy parameters and related weightages to fine tune the RoE results

CAN 5.5 - Advanced RoE and Ticket Matching UI

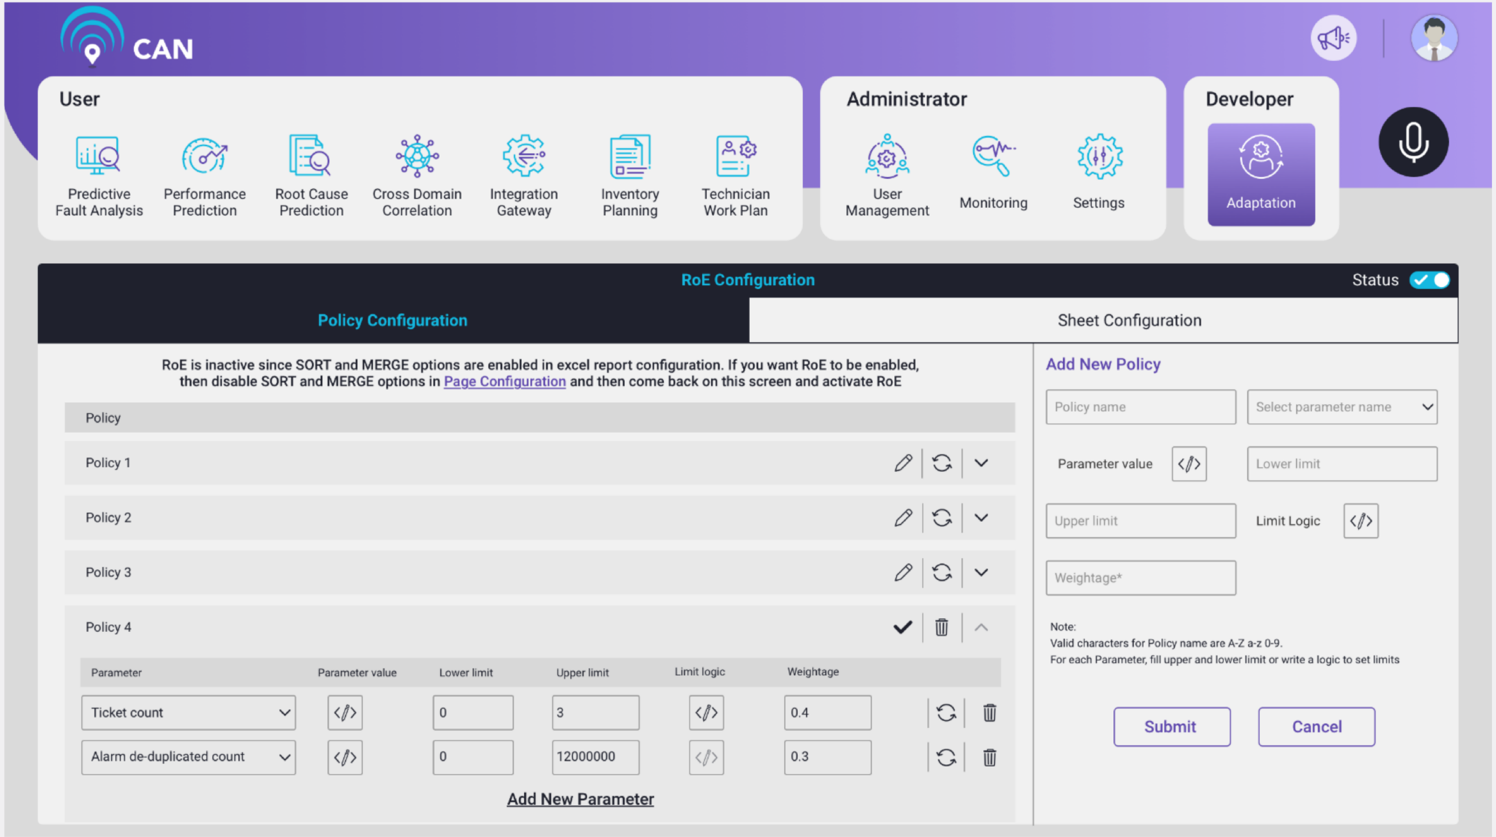

RoE Configuration — Status ✓

Policy Configuration | Sheet Configuration

RoE is inactive since SORT and MERGE options are enabled in excel report configuration. If you want RoE to be enabled, then disable SORT and MERGE options in Page Configuration and then come back on this screen and activate RoE

Policy

- Policy 1
- Policy 2
- Policy 3
- Policy 4

Parameter	Parameter value	Lower limit	Upper limit	Limit logic	Weightage
Ticket count	</>	0	3	</>	0.4
Alarm de-duplicated count	</>	0	12000000	</>	0.3

Add New Parameter

Add New Policy

Policy name	Select parameter name
Parameter value </>	Lower limit
Upper limit	Limit Logic </>
Weightage*	

Note:
Valid characters for Policy name are A-Z a-z 0-9.
For each Parameter, fill upper and lower limit or write a logic to set limits

Submit | Cancel

CAN 5.5 - Tomcat clustering for session management in CAN & CAS

- Clustering of tomcat applications allows a group of tomcat servers serving the incoming HTTP requests instead of one server doing it.
- Enables high availability for the HTTP requests for UI dashboard even if any one of the tomcats in the cluster is down
- Enables session replication or sharing across the clustered tomcats to manage user session
- Session replication enables session recovery in case of server crash as session IDs will be shared across the cluster
- Enables better load management and hence better performance
- Enables better compatibility with kubernetes Horizontal Pod auto-scaling

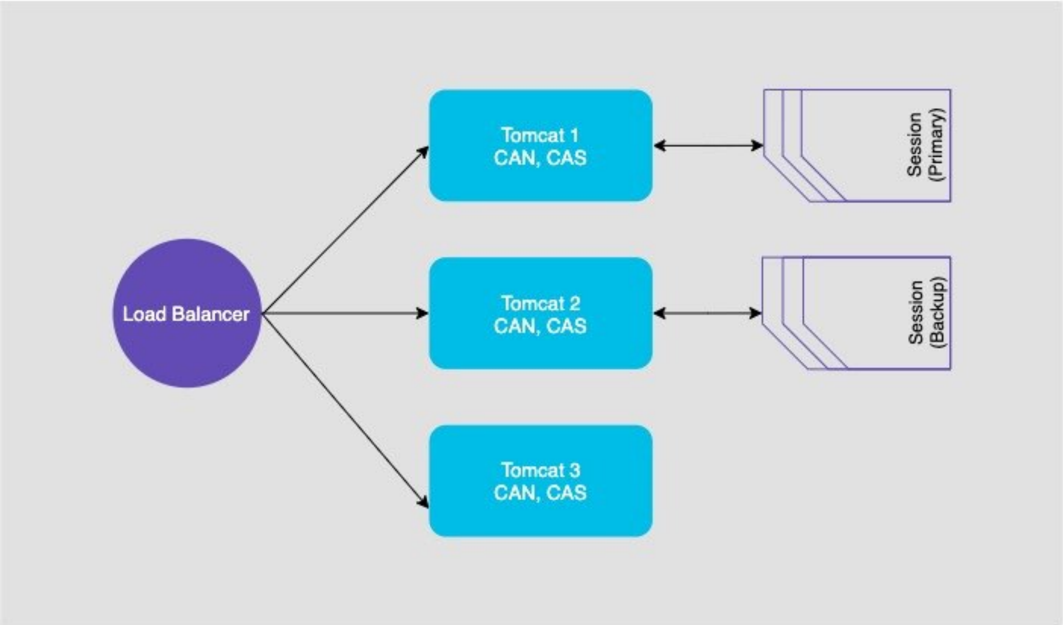

CAN 5.5 - REST API for Prediction Delivery

- REST API for delivering prediction output
- Drops dependency on conventional methods like email
- Secure access through username/password
- User can search and export the required prediction report from CAN database manually or automatically
- Enhance the possibilities of further automation for ticket system integration and other fault management systems
- Enables decoupling of CAN prediction generation and actions reducing the dependency

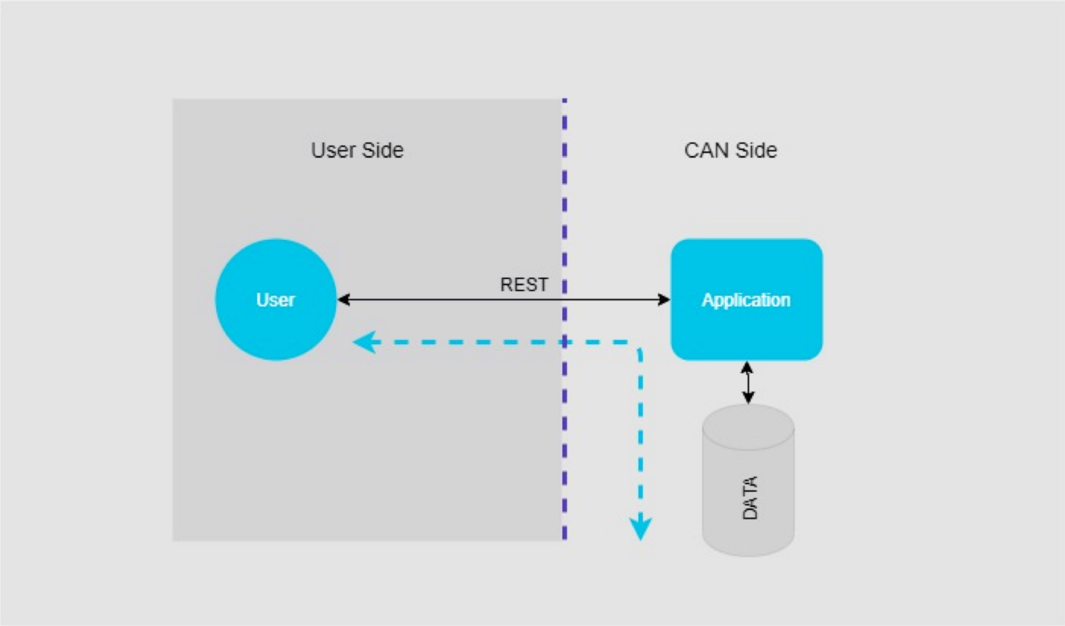

CAN 5.5 - Elasticstack integration in Openshift

- Enables monitoring of the application logs in cloud native environment
- Elastic Stack along with Kibana tool enables the best analysis of application logs in the UI dashboard
- Stores killed pod's container logs for analysis or debugging purpose in case issues

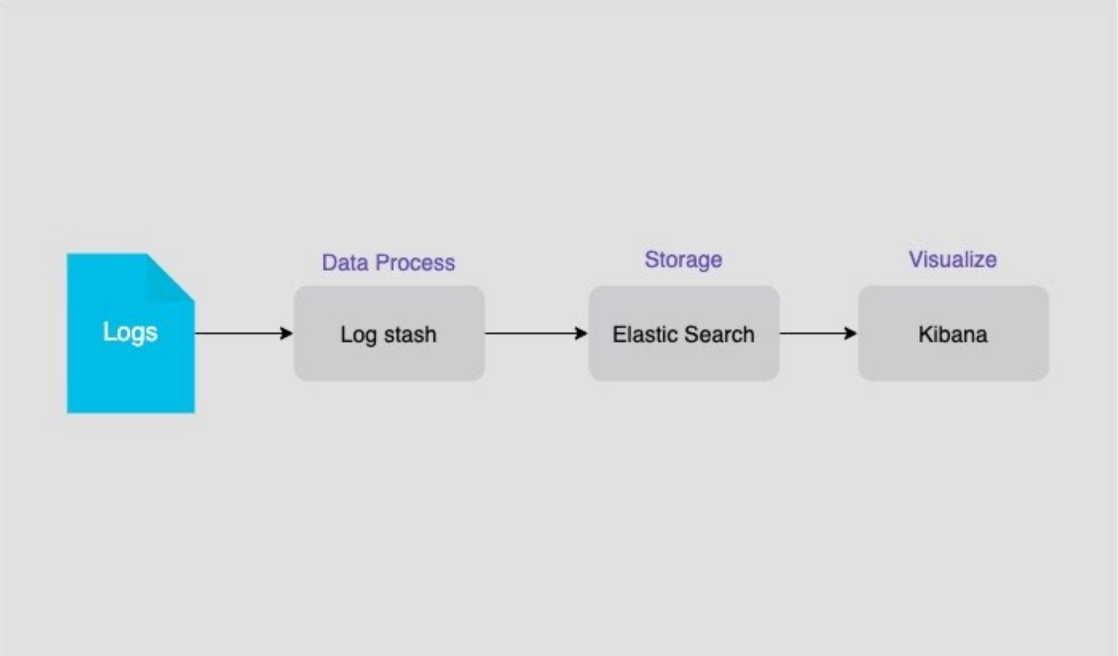

CAN 5.5 - Parser Screen For Work Order Processing

avanseus

- Enhanced parser screen with options like multi selection of rows for addition
- Additional UI Support for Work Order parsing
- Overall, the revamped parser screen enables better user experience for developers/users for input mapping of CAN

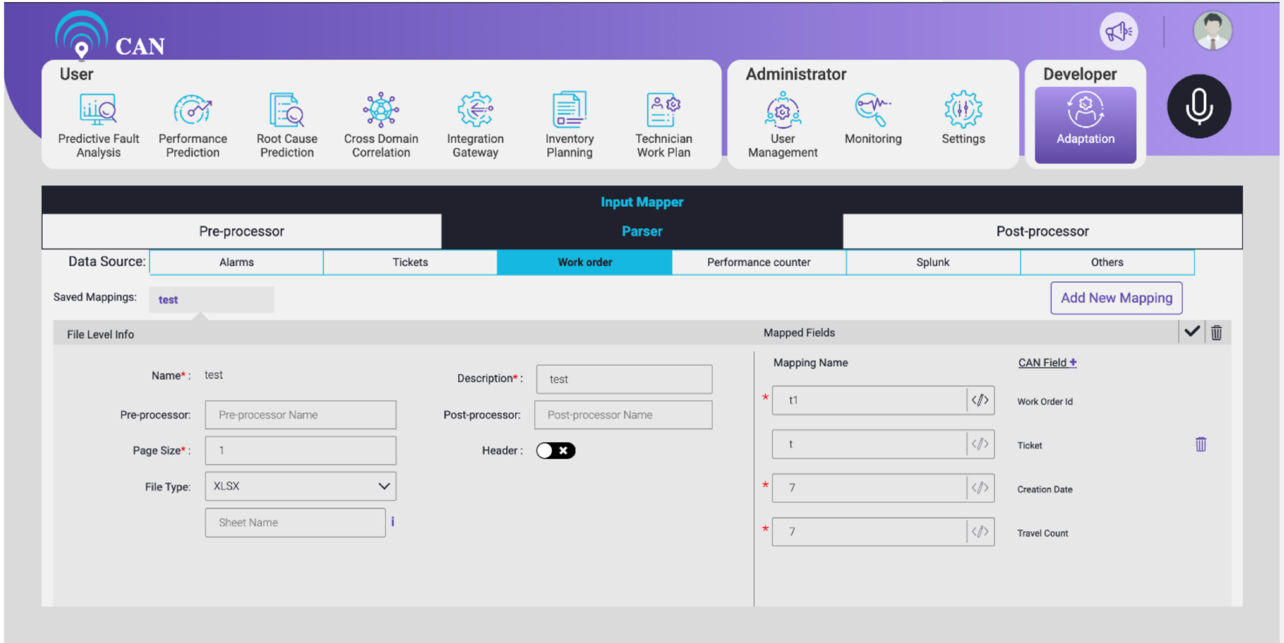

Data Source: Work order

Saved Mappings: test

Name*: test

Description*: test

Pre-processor: Pre-processor Name

Post-processor: Post-processor Name

Page Size*: 1

Header:

File Type: XLSX

Sheet Name

Mapping Name	CAN Field +
* t1	Work Order Id
t	Ticket
* 7	Creation Date
* 7	Travel Count

CAN 5.5 - Horizontal Pod Auto scaling

- Enables auto-scaling of CAN pods based on the CPU Utilization or other customer metrics
- Enables increase or decrease in number of pods based on CPU utilization and enabling optimum utilization of resources and performance
- Highly beneficial when there is an increased load in input requests. In CAN, this applies to increased demand in GUI access by multiple users and batch prediction requests
- Maximum and minimum replica count can be defined for boundary conditions.

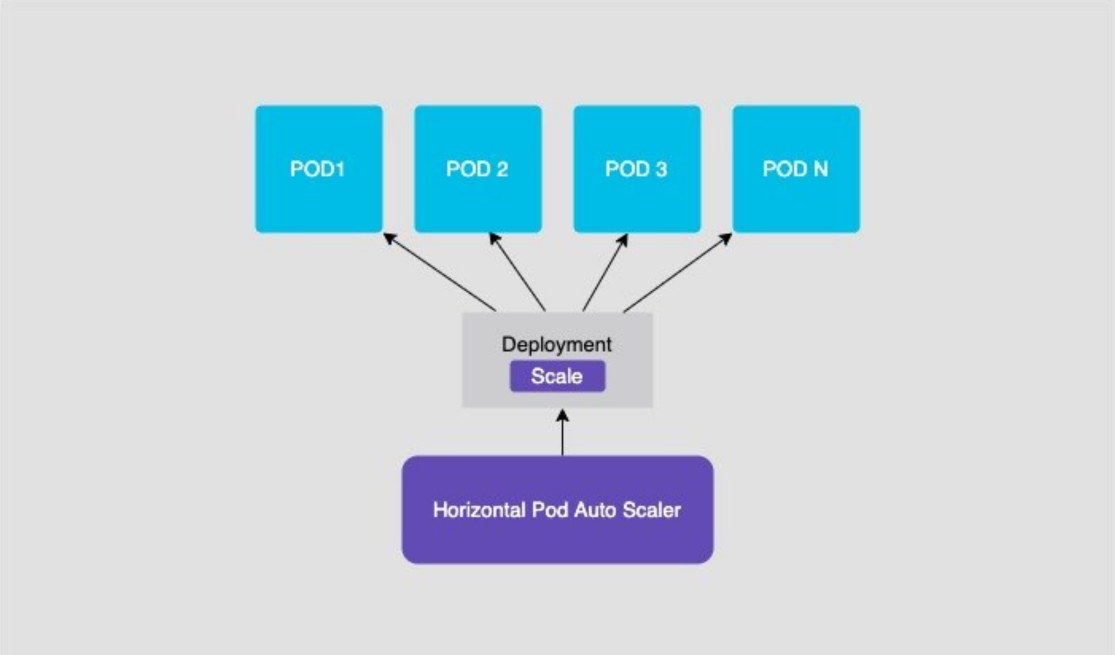

CAN 5.5 - Web security configuration porting from Apache to Nginx

- Porting of all CAN configuration from Apache HTTPD server to Nginix Web Servers
- Improved flexibility in configuration changes
- Easier detection of service affecting activity
- Improved response time
- Higher reliability

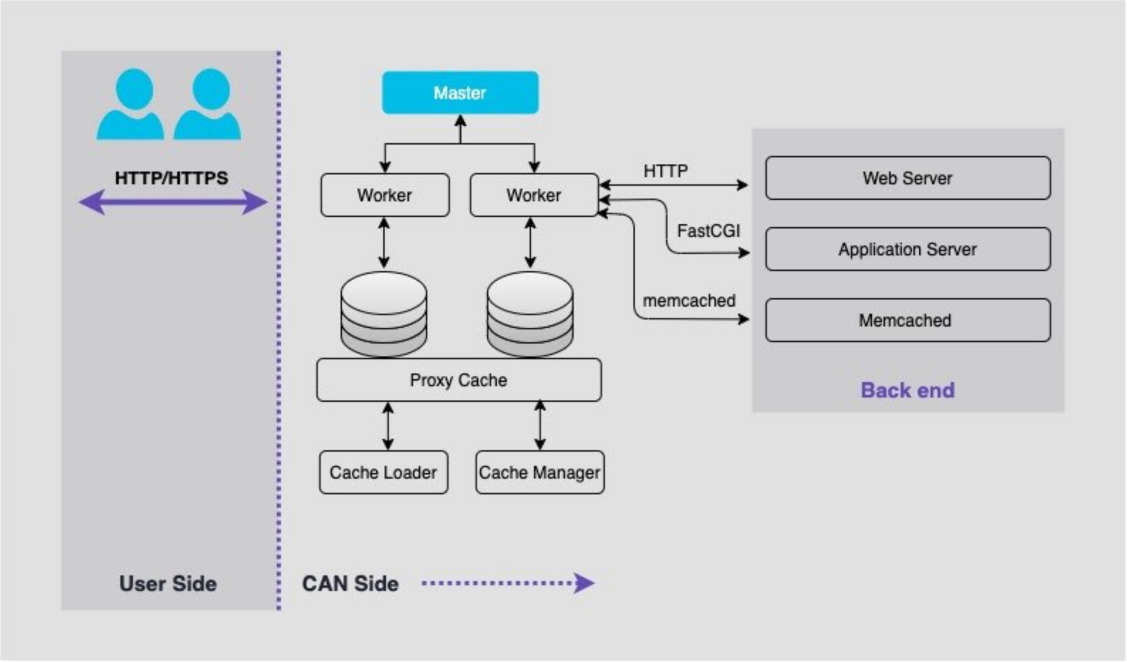

CAN 5.5 - UI Porting with React Framework

- Existing UI is replaced with React JS components which improves response and reliability of UI
- Also separates/decouples the UI modules from backend server module
- Communication between UI & backend modules happens over standard Spring REST API

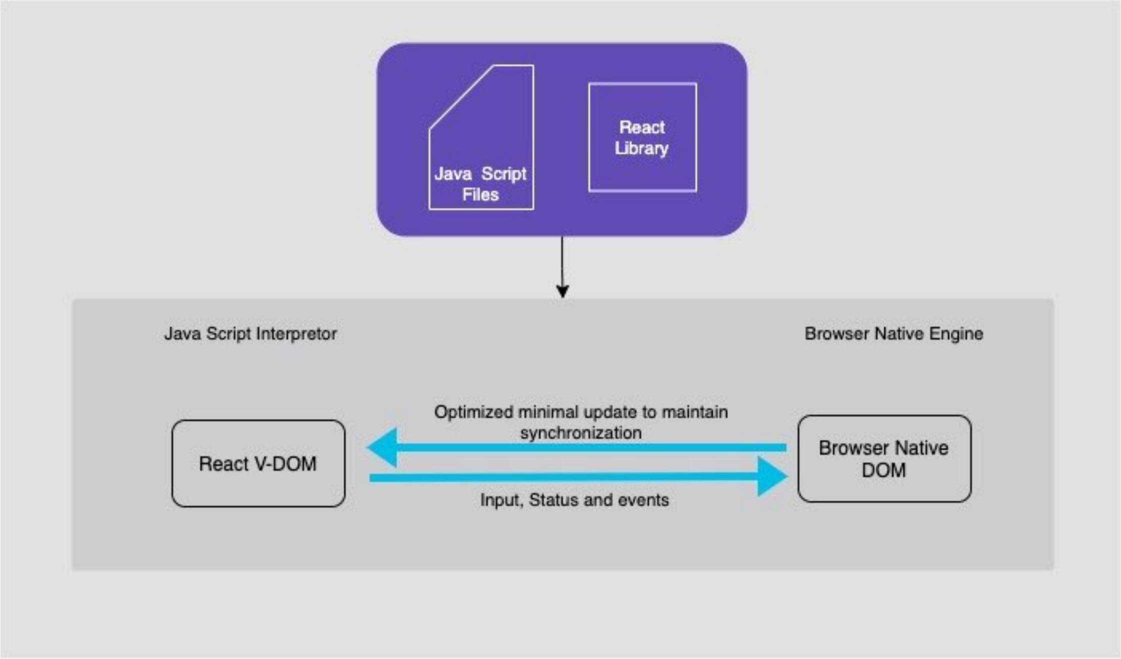

CAN 5.5 - Memcached Tool Integration for Centralized Cache

- Distributed memory caching to store small data
- Caching data storage in centralized server enables faster retrieval
- Improved ease of deployment in high intensity data environment
- Complements the distributed tomcat server setup

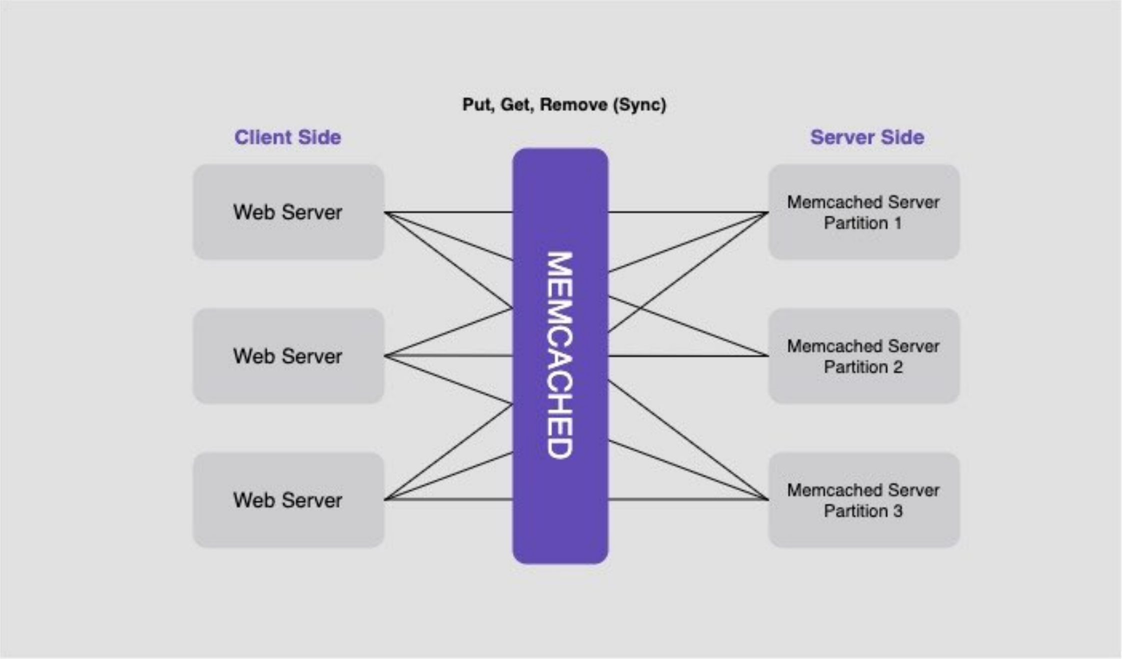
Put, Get, Remove (Sync)
Client Side
Server Side
Web Server
Web Server
Web Server
MEMCACHED
Memcached Server Partition 1
Memcached Server Partition 2
Memcached Server Partition 3

CAN 5.5 - Integration with Istio Service Mesh

- Micro services monitoring for enhanced application efficiency
- Enable automated discovery, load balancing and failure recovery of services
- Added as proxy service enabling Istio layer to remain transparent while retaining the surveillance edge
- Enhanced peer authentication policy and security

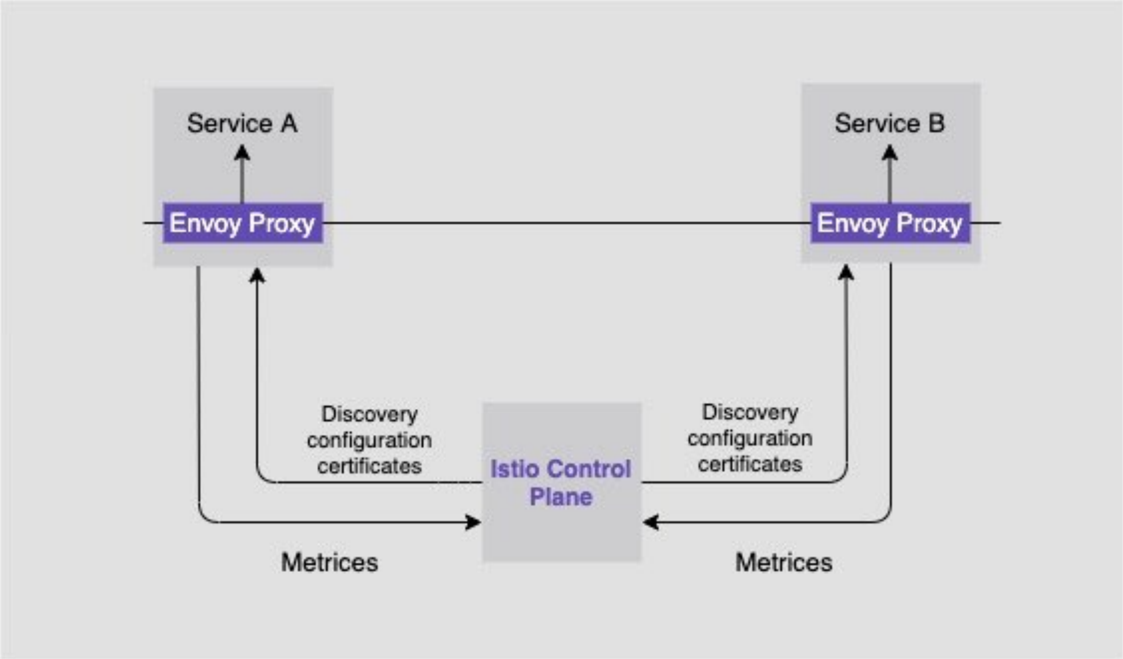

Confidential

CAN 5.5 - Prediction as a Service

- Monolithic prediction process flow converted to modularized architecture
- Decoupling the I/O layer from application layers
- Overcomes the current limitation of 'cause' based division
- Effective load balancing and performance
- Stateless worker nodes enabling quicker processing of predictions

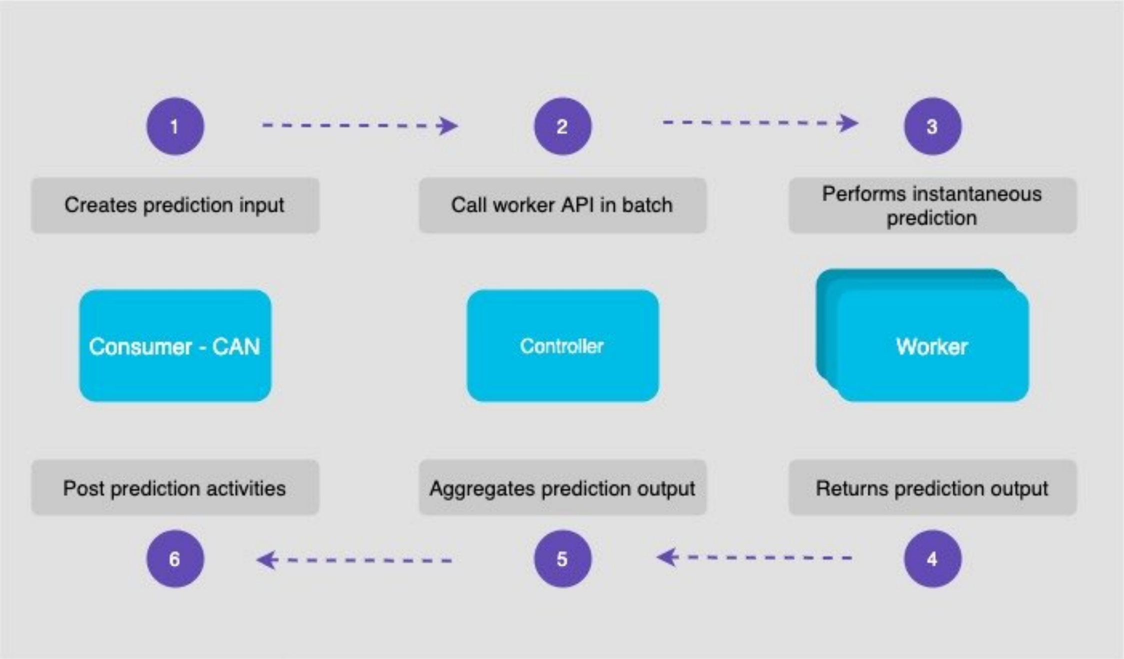
1
2
3
Creates prediction input
Call worker API in batch
Performs instantaneous prediction
Consumer - CAN
Controller
Worker
Post prediction activities
Aggregates prediction output
Returns prediction output
6
5
4

CAN 5.5 - Helm charts for easy installation

- CAN enabled with Helm charts allows enhanced management of cloud native implementations
- Enables easier and quicker management of CAN as regular, upgraded releases and uninstallations can be easily handled through command line interface.

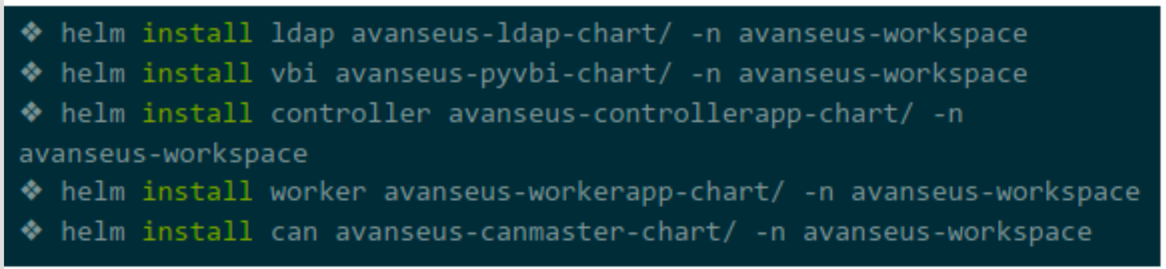

```
❖ helm install ldap avanseus-ldap-chart/ -n avanseus-workspace
❖ helm install vbi avanseus-pyvbi-chart/ -n avanseus-workspace
❖ helm install controller avanseus-controllerapp-chart/ -n
avanseus-workspace
❖ helm install worker avanseus-workerapp-chart/ -n avanseus-workspace
❖ helm install can avanseus-canmaster-chart/ -n avanseus-workspace
```

Single Command Line for CAN installation through Helm

CAN 5.5 - Integration with Monitoring Softwares (Kiali, Promotheus and Grafana)

- Enhanced monitoring of CAN Services by monitoring traffic flow to infer the topology and identify errors using Kiali management consoles.
- Realtime display of Service Mesh health based on the traffic topology
- All events within the service mesh to be closely monitored by Prometheus for any events or faults and records real time metrics in time series
- These identified metrics and data are displayed graphically using Grafana as “Avanseus Dashboard”
- Continuous monitoring of inbound and outbound traffic of every interested micro services, memory, and CPU utilization of all the micro services, stats related to each micro services, etc

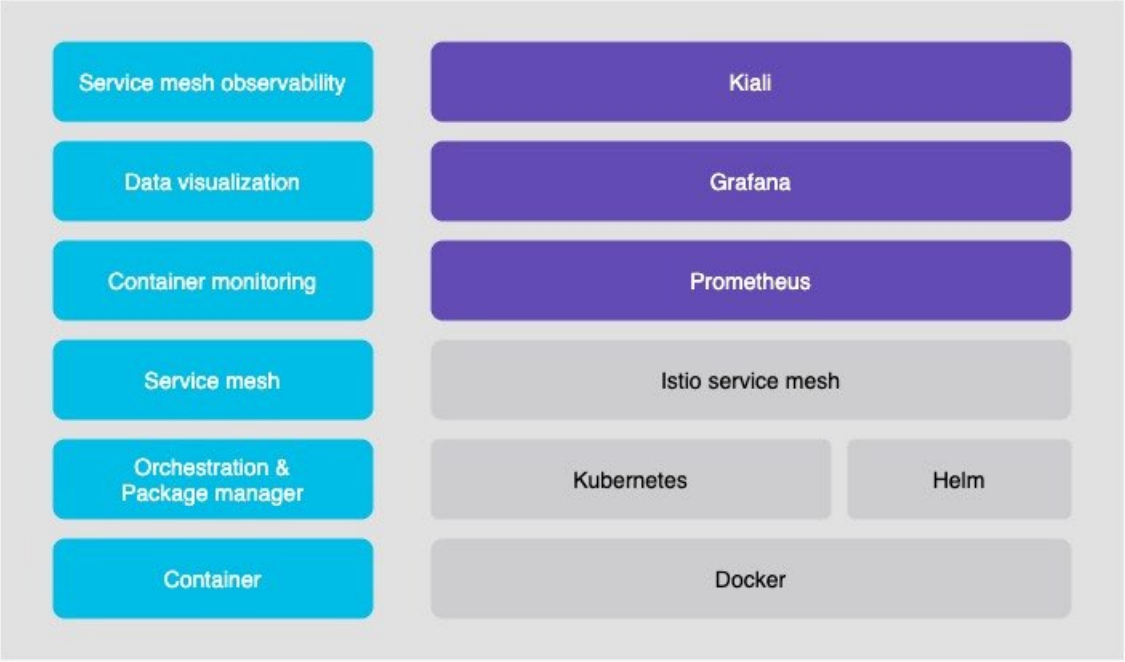

Service mesh observability	Kiali	
Data visualization	Grafana	
Container monitoring	Prometheus	
Service mesh	Istio service mesh	
Orchestration & Package manager	Kubernetes	Helm
Container	Docker	

CAN 5.5 - Re-scheduling of trigger based on UI cron pattern configurations

- UI Support for scheduling CAN activities like input data processing, prediction generation etc
- Enabling more flexibility to the users to schedule the activity with zero coding
- Enhanced user friendliness

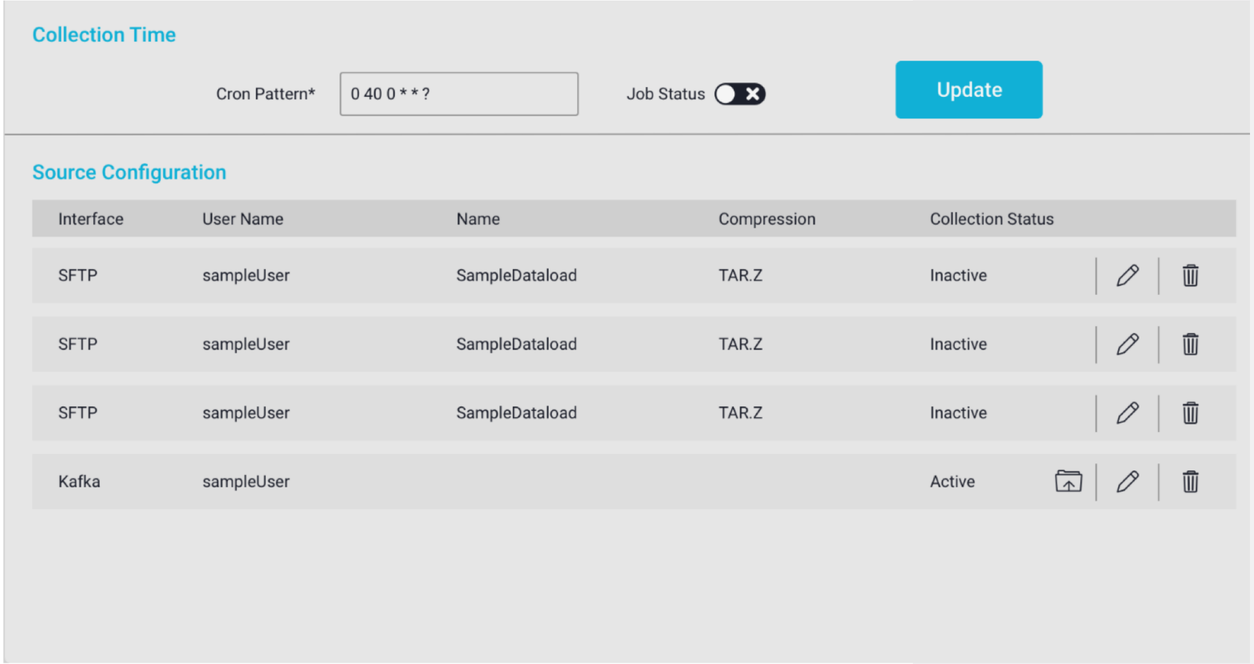

Collection Time

Cron Pattern* 0 40 0 * * ?

Job Status

Source Configuration

Interface	User Name	Name	Compression	Collection Status
SFTP	sampleUser	SampleDataload	TAR.Z	Inactive
SFTP	sampleUser	SampleDataload	TAR.Z	Inactive
SFTP	sampleUser	SampleDataload	TAR.Z	Inactive
Kafka	sampleUser			Active