

# Juniper Automation Awareness Anuta ATOM Bootcamp (Day-4)

Kiran Sirupa

Head of  
Marketing

Vinay US

DevOps Lead

Sitaram Jakka

Sr. Solution  
Architect

Ashok  
Tippireddy

Engineering  
Manager



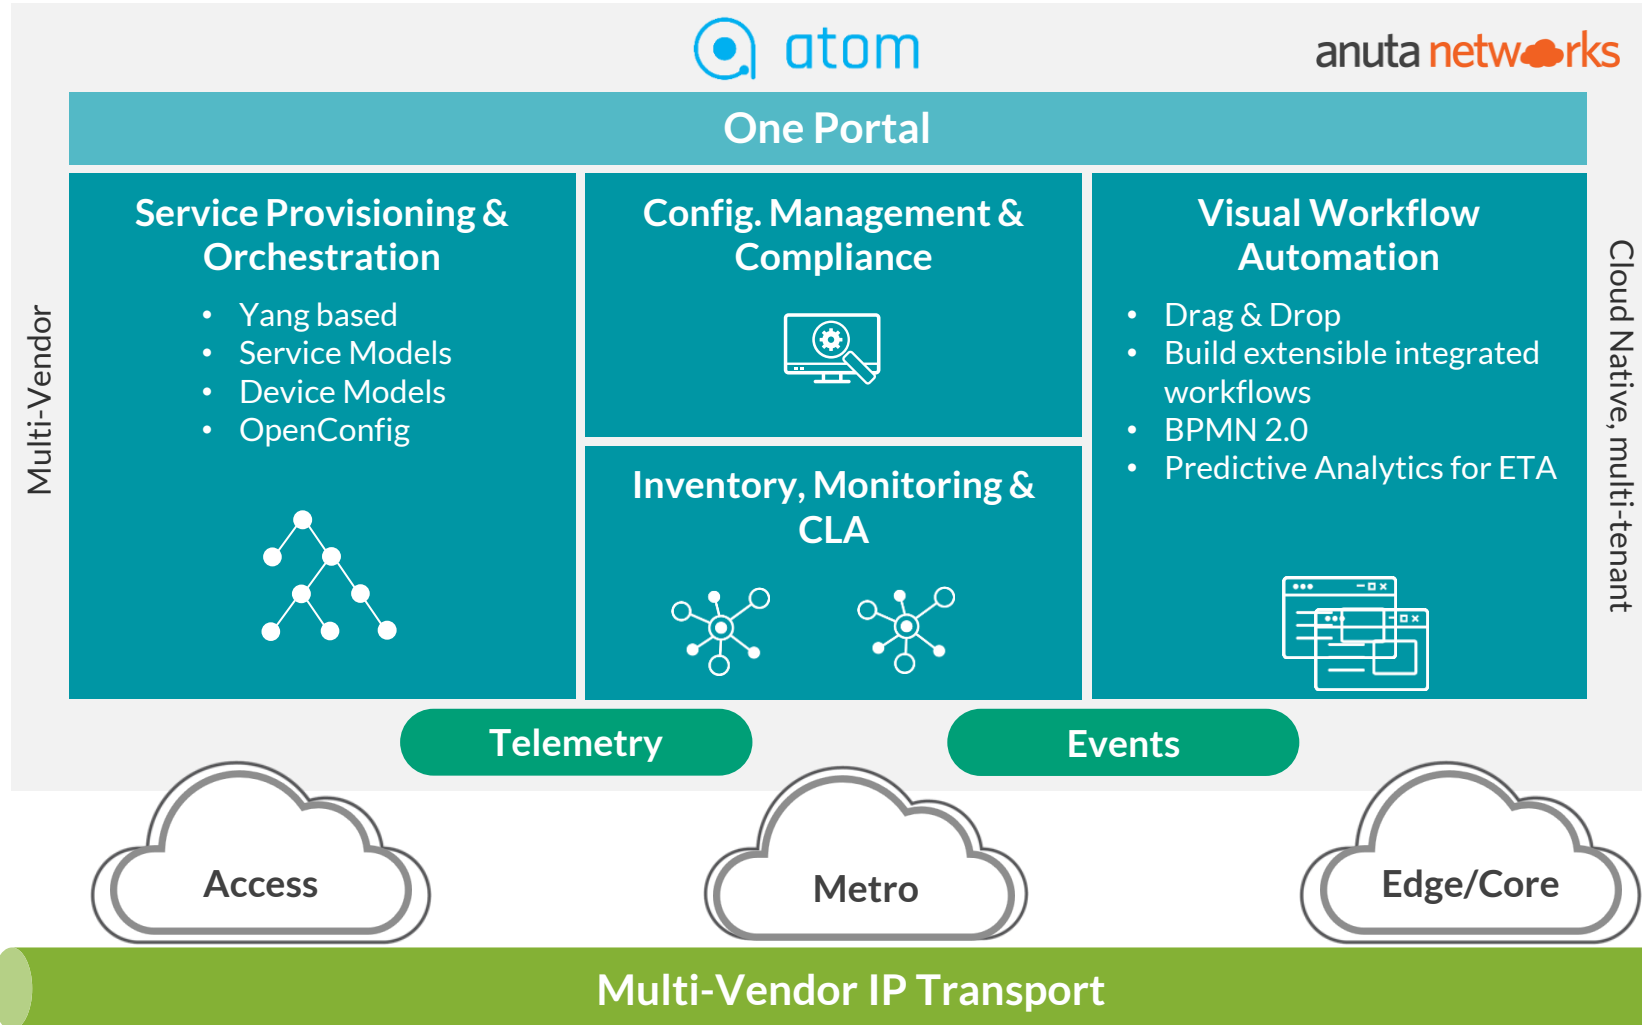
# AGENDA

---

- Day – 1: Compliance Management
- Day – 2: Workflow Hands-On
- Day – 3: Service Model Hands-On
- Day – 4: ATOM Deployments
- Administration & Troubleshooting
- Customer Deployments Deep-Dive
- Closed-Loop Automation
- Scale, HA, Performance
- Licensing

# Anuta ATOM

## Service orchestration, Workflow and Compliance with Closed loop Automation



- 3<sup>rd</sup> Party Resale agreement
- On Juniper pricelist
- JTAC
- Juniper Pro services

# ATOM – Administration and Troubleshooting

# Agenda

## Dashboards

- System Manager
- Kibana
- Grafana
- Prometheus
- Keycloak
- Glowroot\*

## CLI

- kubectl
- helm
- k9s\*

# System Manager

- Inhouse management portal
- Accessed within the ATOM GUI
- Enable/Disable maintenance mode
- View component relations
- View component functions
- View component application
- View deployments

# Kibana

- Logging management portal for ELK stack
- Access via the context path `/kibana` for the GUI
- View logs of tasks and events
- Data is captured via log-forwarder daemonset into elasticsearch cluster
- Provides filter support
- Default indexes contains application and event data
- Custom indexes can be created for analytics
- Custom dashboards can be created for visualization

# Grafana

- Monitoring portal for time series data
- Access via the context path **/grafana** for the GUI
- View visual representation of health and events
- Good view of node stats and kafka events
- Data is captured via prometheus
- Provides filter support
- Custom dashboards can be created for visualization



# Prometheus

- Metrics collector module
- Access via the context path **/prometheus-atom** for the GUI
- GUI can be used to run queries for events and trends
- Rules are used to configure alerts
- Provides dashboard view as well
- Custom rules and alert notifications can be invoked
- Supports webhook support such as slack

# Keycloak

- Authentication module
- Access via the context path `/auth` for the GUI
- GUI can be used to configure authentication information
- SSO is provided here
- Redirect URLs and additional contexts support
- Multiple clients on a tenant support
- Multi tenancy handling via Realm

# Glowroot(optional)

- JVM stat collector module
- Access via the context path **/glowroot** for the GUI
- GUI can be used to view JVM stats of each microservice
- Memory, CPU and GC info can be visualised at a deeper level

# kubectl

- Native kubernetes command line tool
- Needs access to control plane node or kubernetes config file for remote access such as EKS
- Interacts with kubernetes api-server
- Can be used to manually check health, logs and events

Sample commands:

```
kubectl get nodes
```

```
kubectl describe pod -n <namespace> <pod name>
```

```
kubectl logs -n <namespace> <pod name>
```

# helm

- Package manager for deploying kubernetes objects
- ATOM uses helm to deploy its microservices
- Supports upgrade, rollback and history management
- Interacts with kubernetes api-server

Sample commands:

```
helm ls -n <namespace>
```

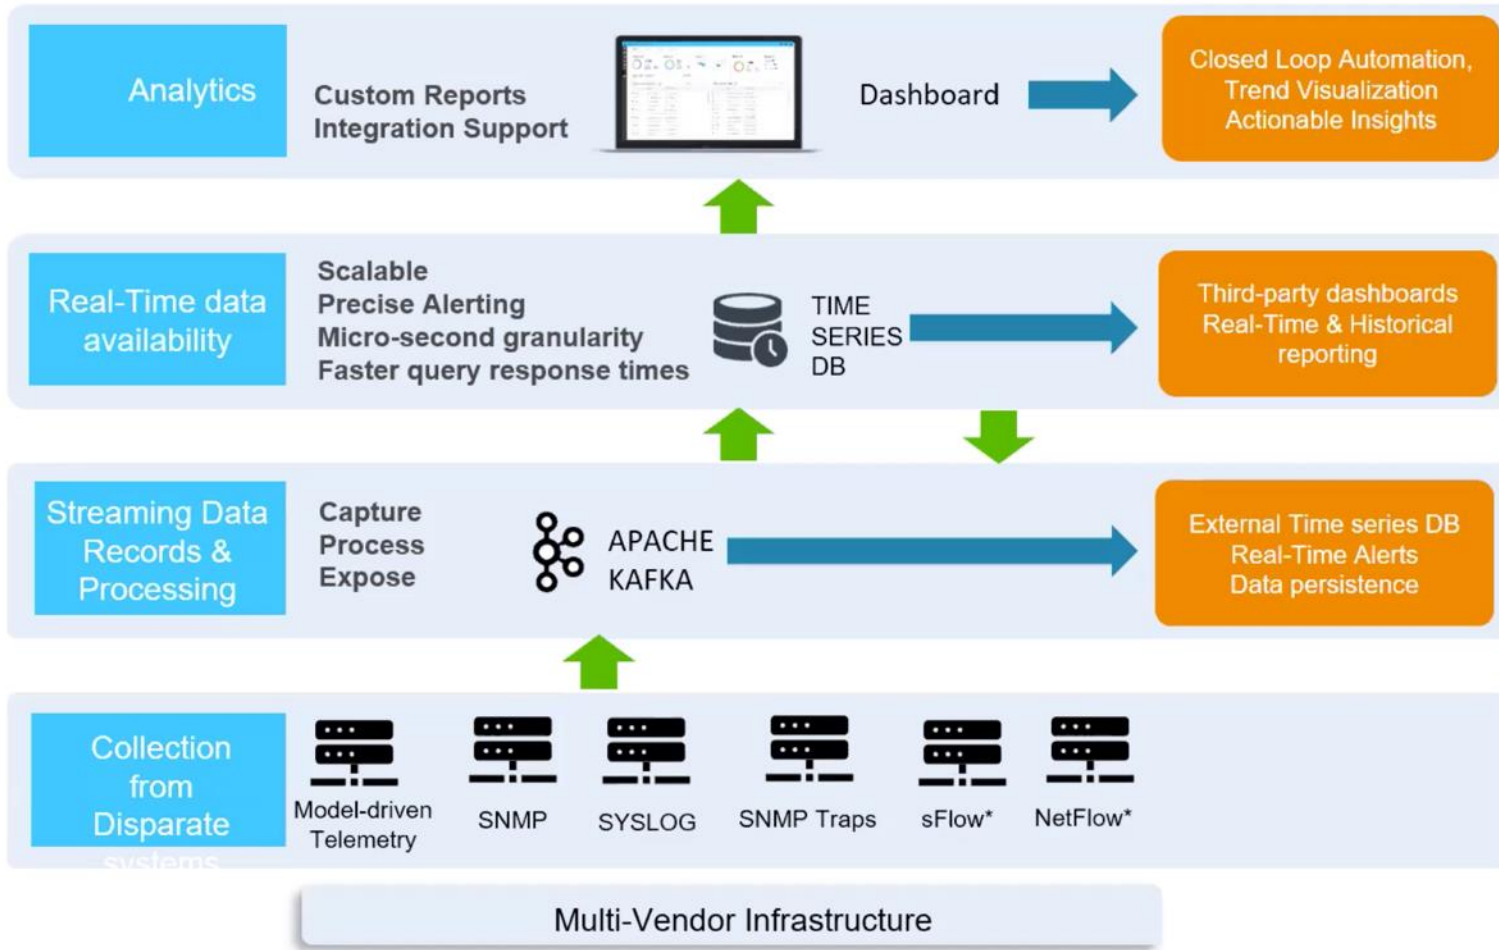
```
helm install <chart name> <path or tgz file> -n <namespace>
```

```
helm history <chart name> -n <namespace>
```

# k9s(optional)

- Open Source kubernetes task manager invoked at CLI
- Interacts with kubernetes api-server
- Avoids typing commands through kubectl
- Can be used to run all kubernetes operation via its CLI dashboard
- Can be used to view, update and delete any object

# Telemetry and Real-time Analytics



Monitor with single pane of glass

Minimize MTTR

Integrate with external databases

Customizable dashboard

# ATOM Monitoring capabilities

## Support for Legacy protocols

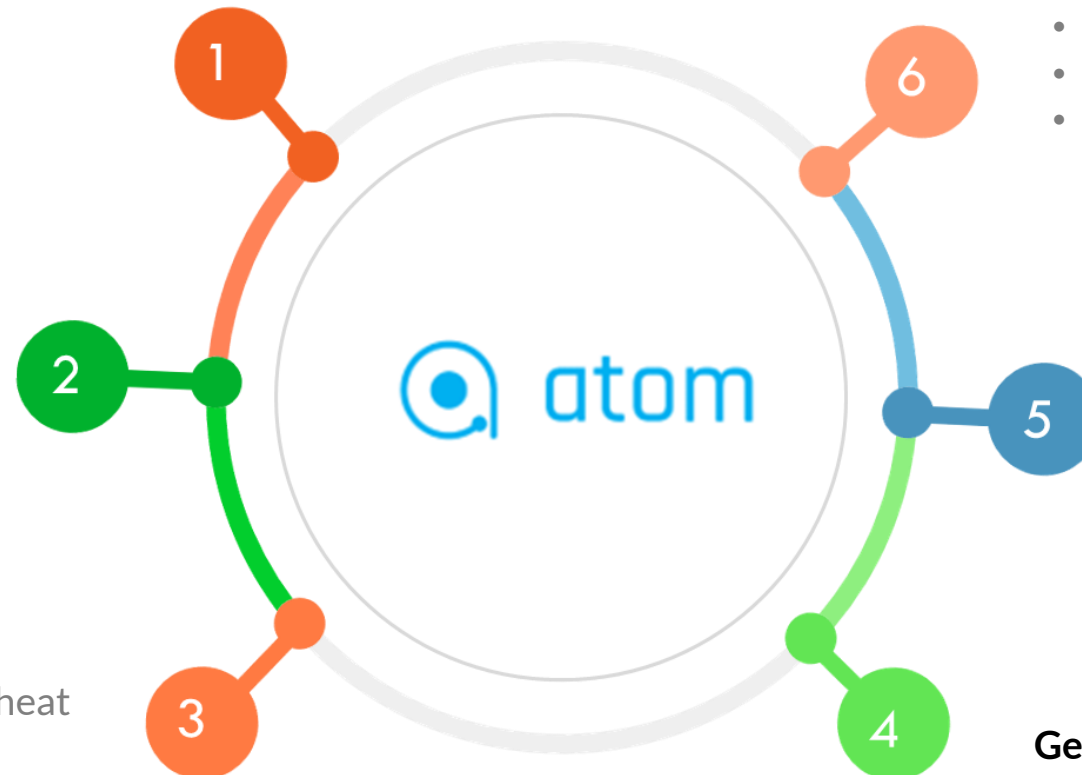
- SNMP
- SNMP Traps
- Syslog

## Streaming Telemetry

- Transport - gRPC, TCP
- Dial modes and packet encoding
- OpenConfig & Native model

## In-depth KPI monitoring

- Platform level statistics & heat maps
- Routing protocol metrics
- In-depth interface statistics
- Customizable Charts



## Time-series database

- Horizontally scalable
- High precision data collection
- Millions of metrics
- Roll-up & Down sampling

## Multi-Vendor Support

- 45+ Vendors
- 100+ platforms

## Geo-Distributed Architecture

- Microservices architecture
- Aids low-latency requirements



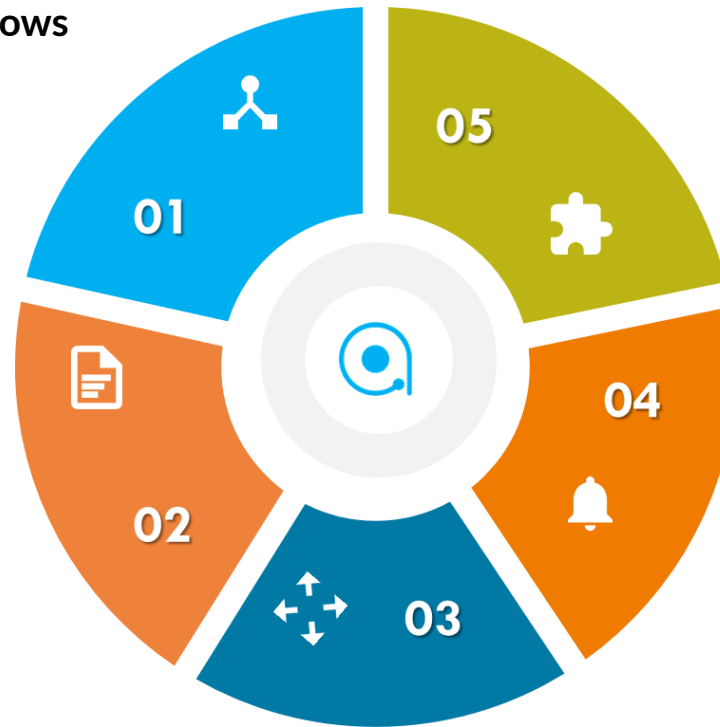
# ATOM Alerting capabilities

## Trigger Advanced Workflows

- Automate Remediation flows
- Event Enrichment
- Event Correlation

## Policy Definitions

- Mix & Match conditions
- Looping
- Windowing



## Auto-ticketing

- ServiceNow
- Jira

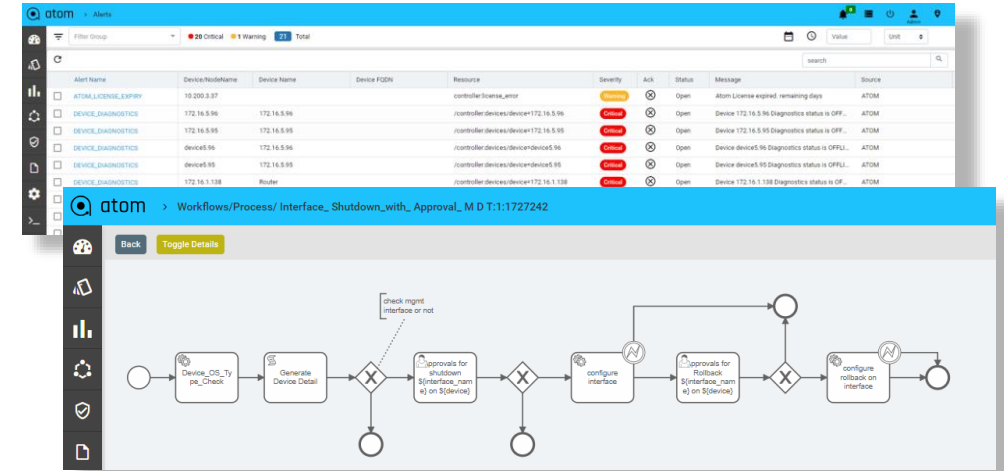
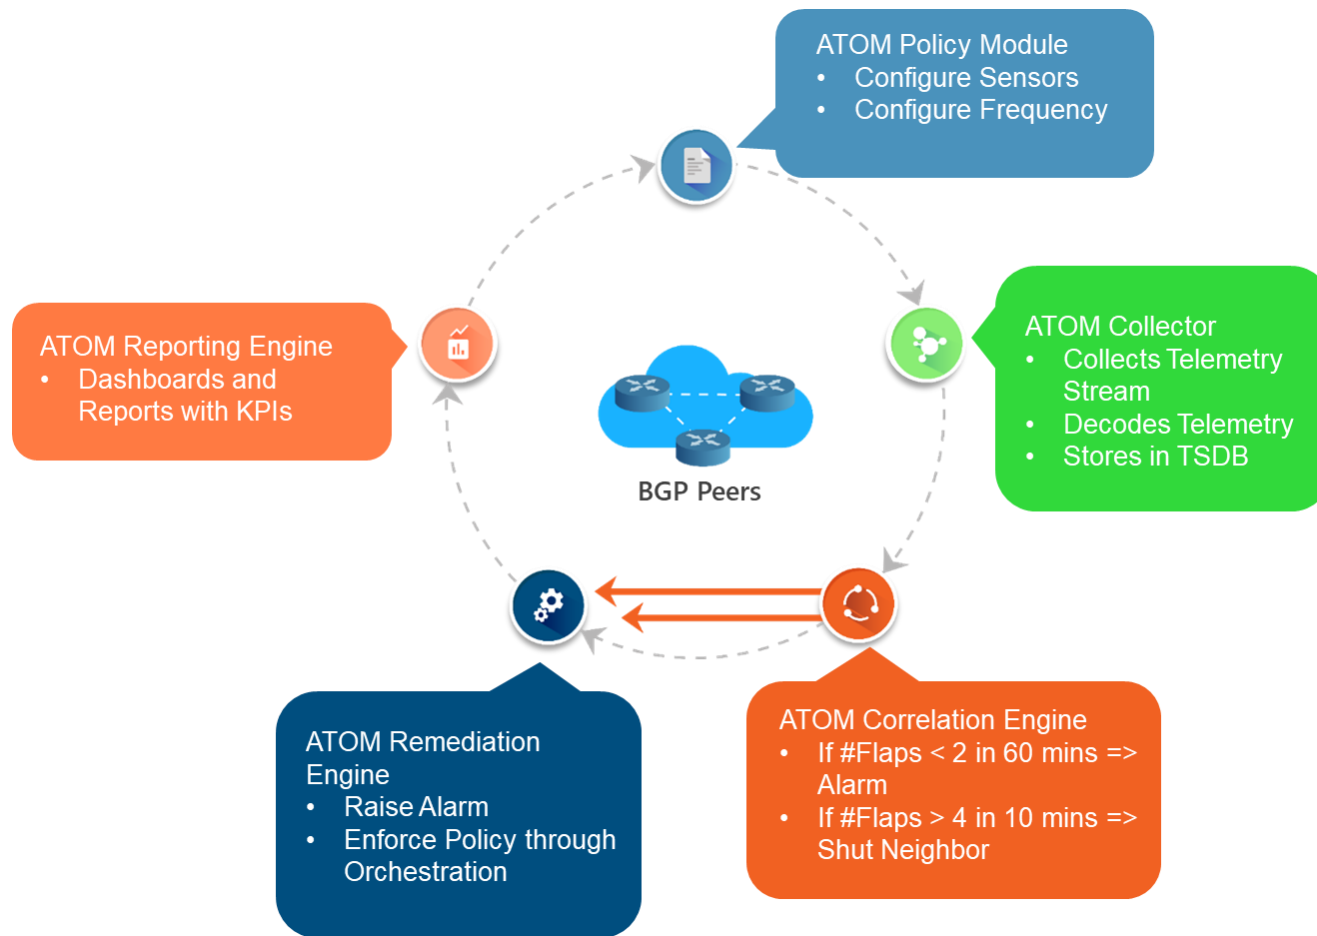
## Automated Alerting & Alert Routing

- Email
- Slack
- Alarm
- Webhooks

## Open Platform

- External collectors
- Event enrichment DBs

# Closed Loop Automation – BGP Neighbor Flapping



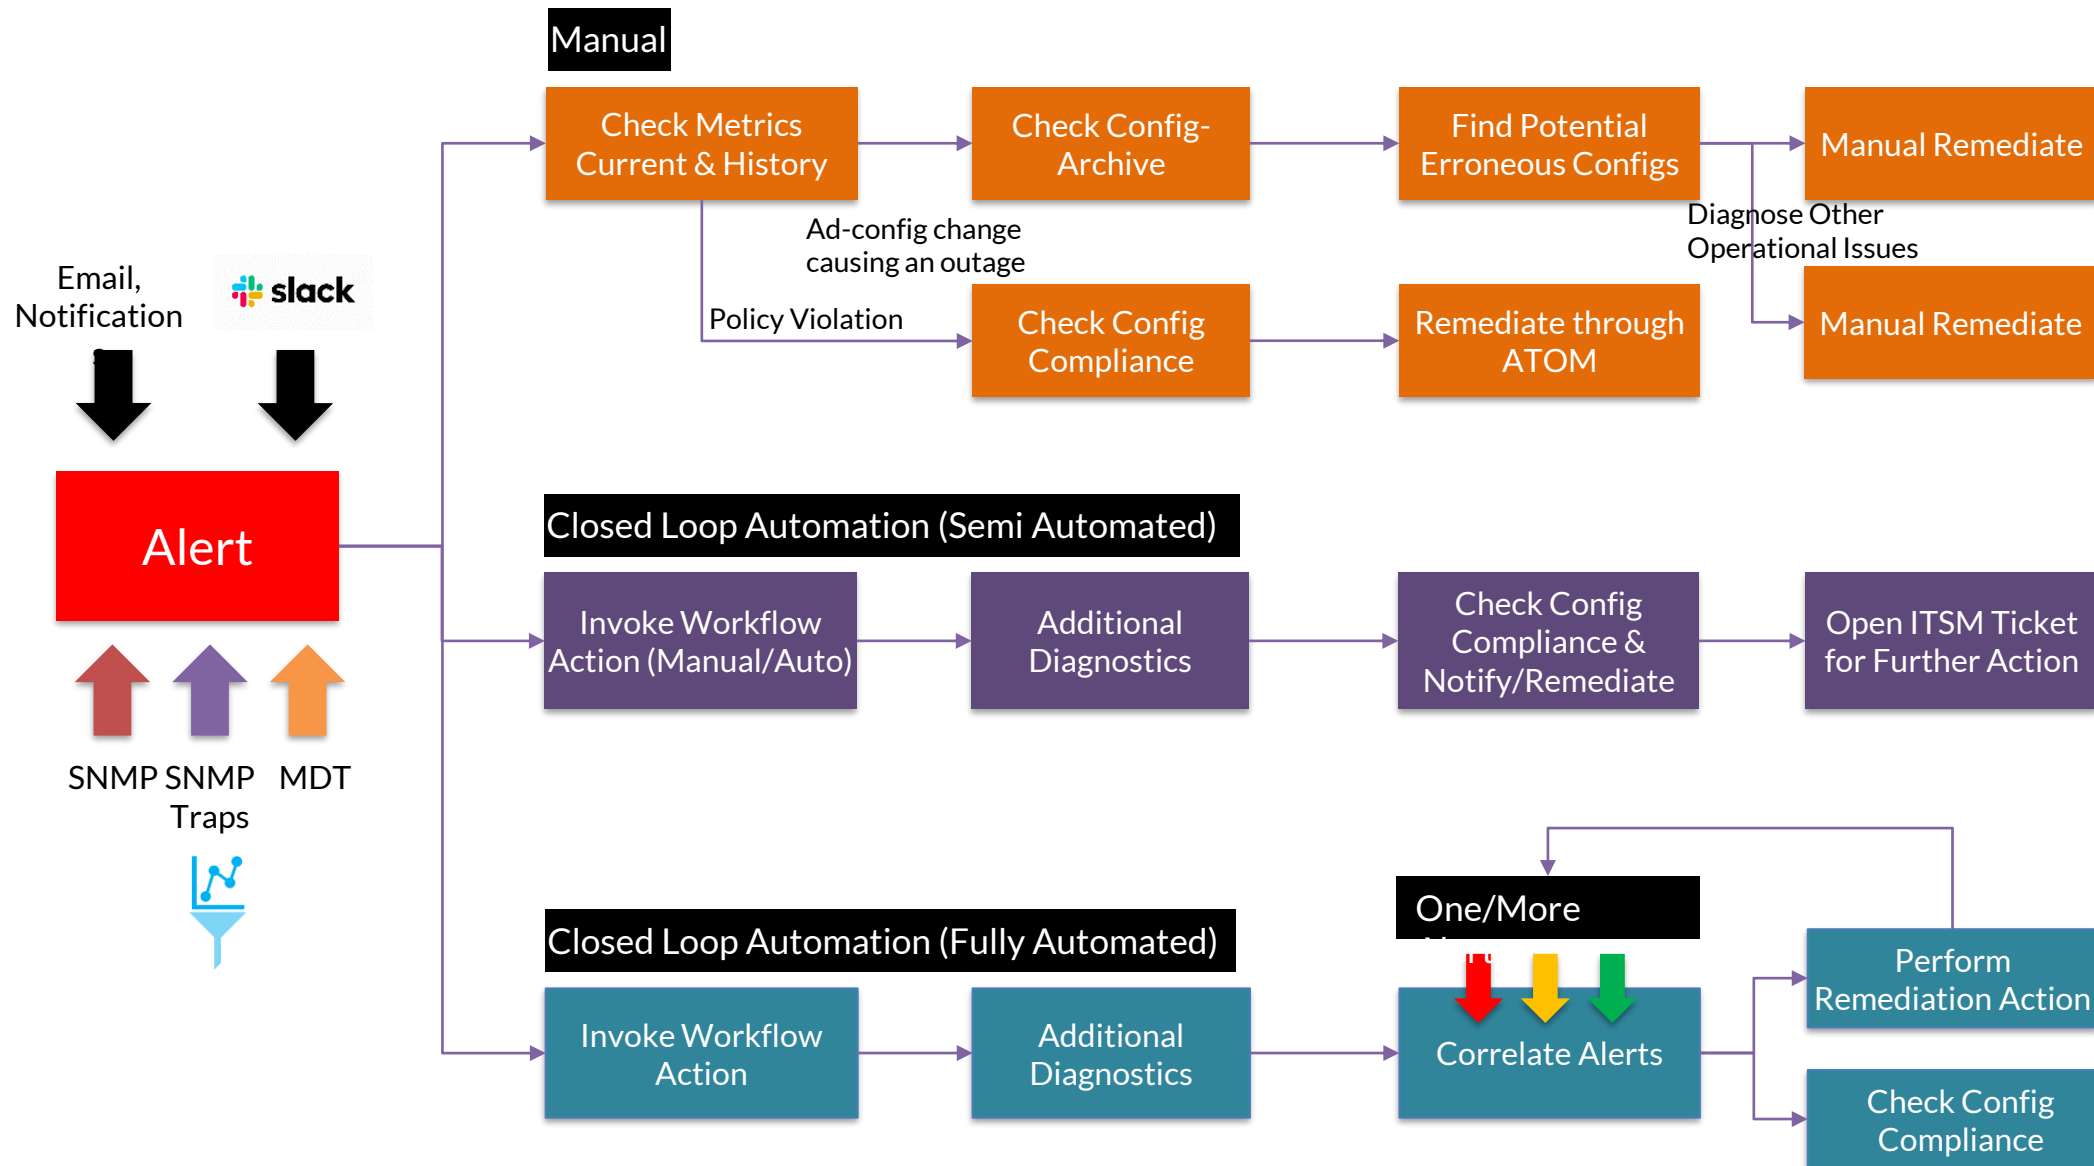
Meet Your Customer SLAs

Eliminate Downtimes

Reduce MTTR, Increase Efficiency

Reduce OPEX

# Anuta ATOM enables Closed Loop Automation



# Anuta ATOM supports several prebuilt templates

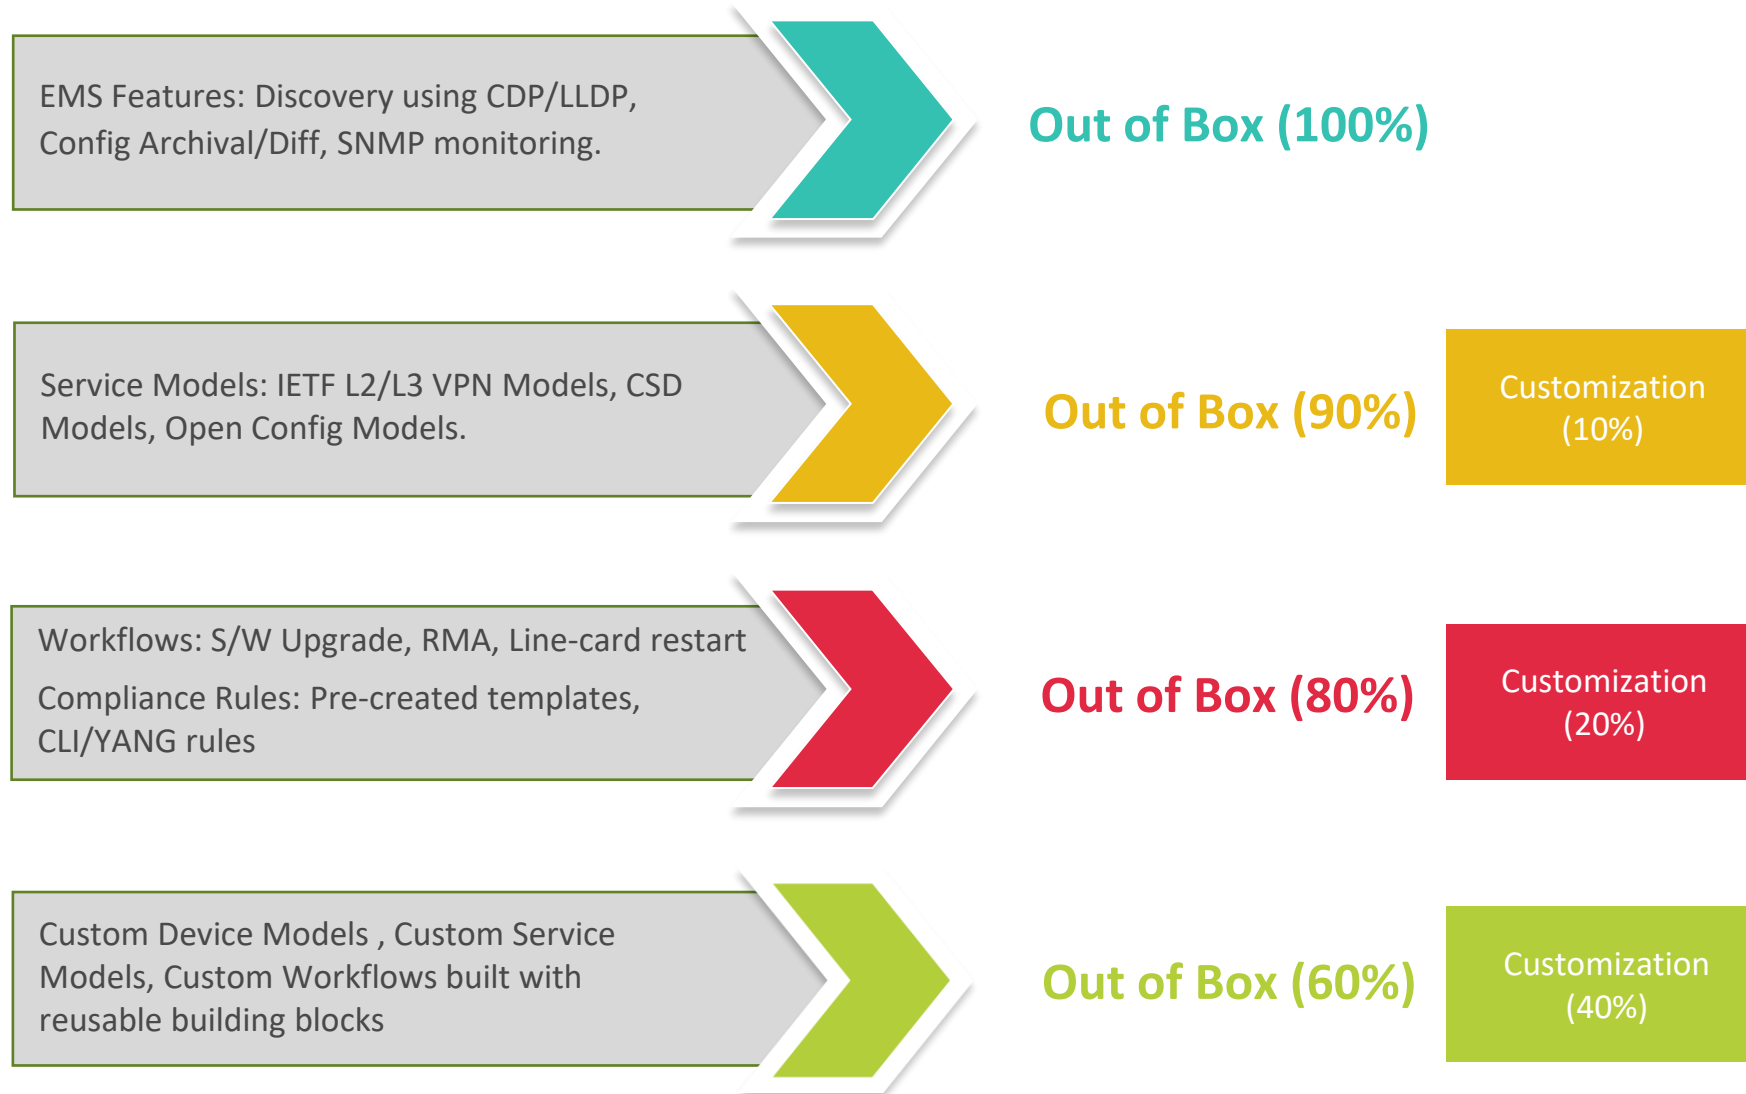

Shipping

Committed

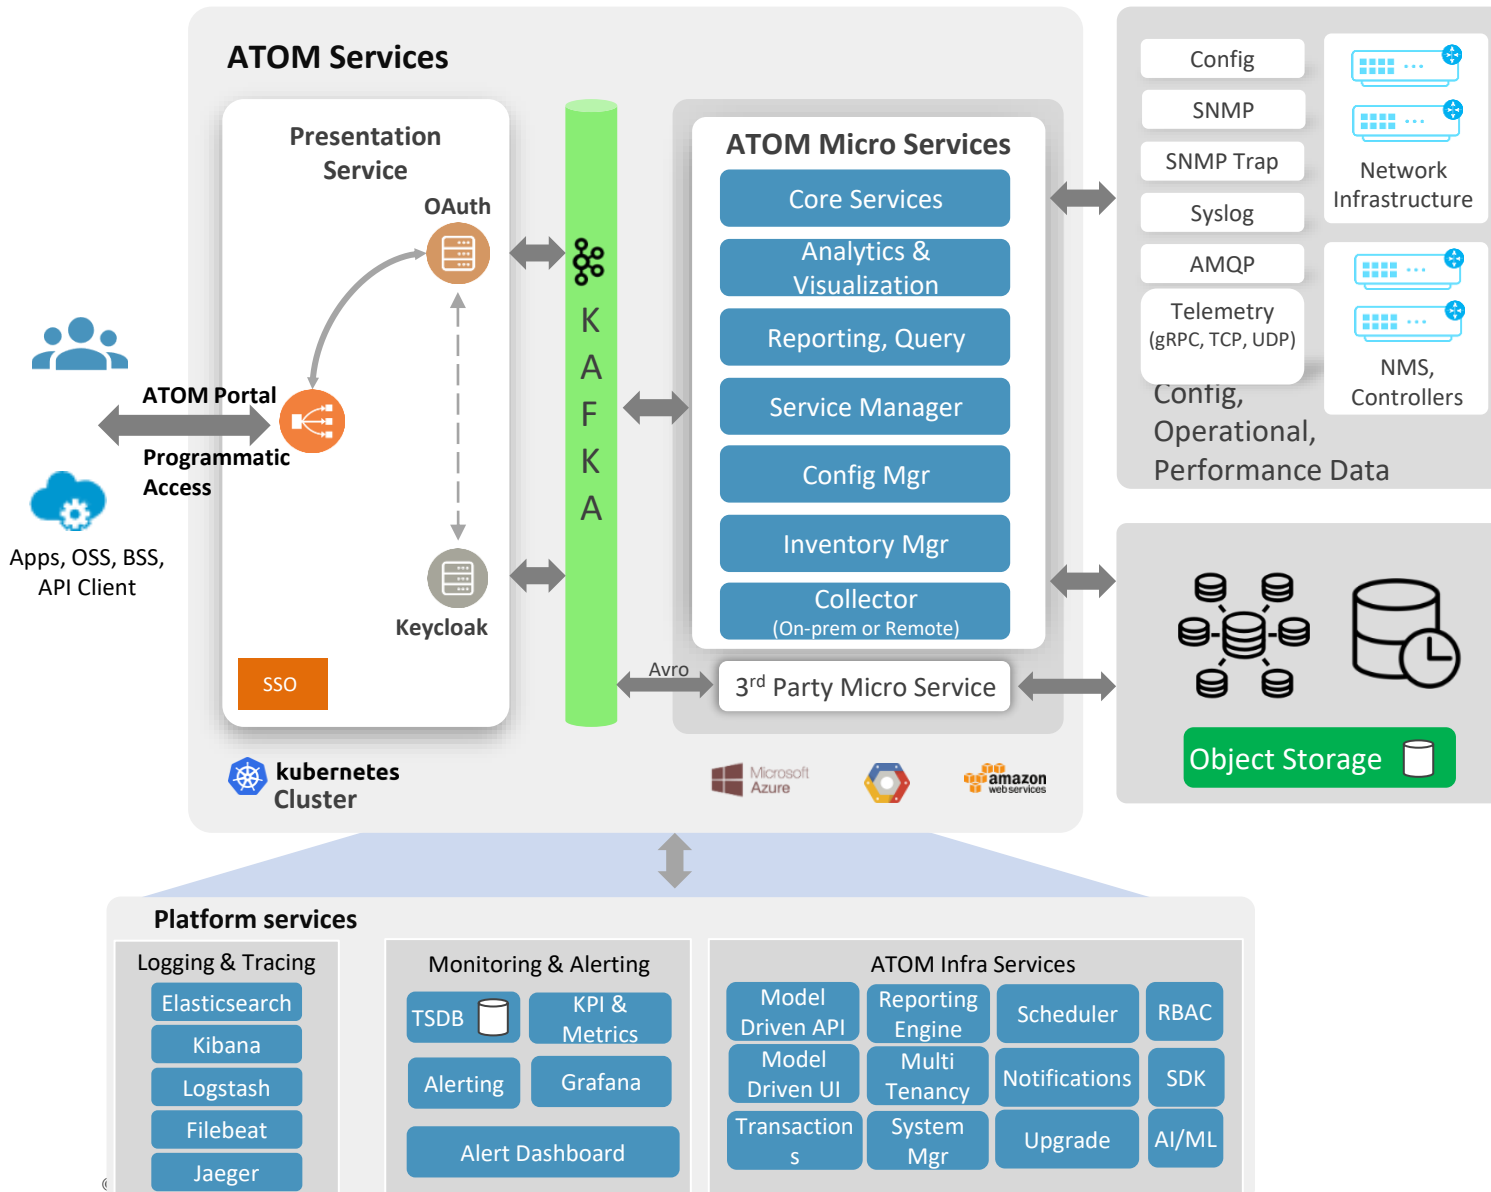
Q2-2021	Q3-2021	Q4-2021
<ul style="list-style-type: none"> <li>✓ ELine-Dot1q-SingleVLAN</li> <li>✓ ELine-PortBased</li> <li>✓ ELine-QinQ-AllVLAN</li> <li>✓ ELine-QinQ-VLANRange-CCC</li> <li>✓ ELine-BGP-Port-Based</li> <li>✓ ELAN-BGP-Dot1Q-SingleVLAN</li> <li>✓ ELAN-BGP-PortBased</li> <li>✓ ELAN-BGP-QinQ-AllVLAN</li> <li>✓ ELAN-Hub-Spoke-QinQ-AllVLAN</li> <li>✓ L3VPN-OSPF-STATIC L3 VPN (Full Mesh)</li> <li>✓ L3VPN-BGP-STATIC L3 VPN (Full Mesh)</li> <li>✓ L3VPN-OSPF-Static (Hub-Spoke-1-Interface)</li> <li>✓ L3VPN-BGP-Static (Hub-Spoke-1-Interface)</li> <li>✓ EVPN-VXLAN</li> <li>✓ EVPN-ETREE</li> </ul>	<ul style="list-style-type: none"> <li>✓ ELine-Dot1q-SingleVLAN-CCC</li> <li>✓ ELine-Dot1q-SingleVLAN-Ext-CCC</li> <li>✓ ELine-QinQ-AllVLAN-CCC</li> <li>✓ ELine-QinQ-AllVLAN-Ext-CCC</li> <li>✓ ELine-QinQ-VLANRange</li> <li>✓ ELine-QinQ-VLANRange-Ext-CCC</li> <li>✓ Eline-BGP-QinQ-AllVLAN</li> <li>✓ Eline-BGP-Dot1q-SingleVLAN</li> <li>✓ ELAN-BGP-Dot1q-Normalized-VLAN-None</li> <li>✓ ELAN-BGP-QinQ-AllVLAN-Normalized-All</li> <li>✓ ELAN-BGP-Dot1q-Normalized-VLAN-None</li> <li>✓ ELAN-BGP-QinQ-Range-Normalized-VLAN</li> <li>✓ ELAN-Hub-Spoke-QinQ-AllVLAN-No</li> </ul>	<ul style="list-style-type: none"> <li>✓ Inventory Management &amp; Reporting (parity with Junos Space)</li> <li>✓ Service Functional Audit</li> <li>✓ Service Monitoring Dashboards</li> <li>✓ Service Verification</li> </ul>

**Note:** Junos Space Connectivity Services Director (CSD) is EoS and the next-generation product for CSD is Anuta ATOM

# Anuta ATOM – Out of Box Support and Customization

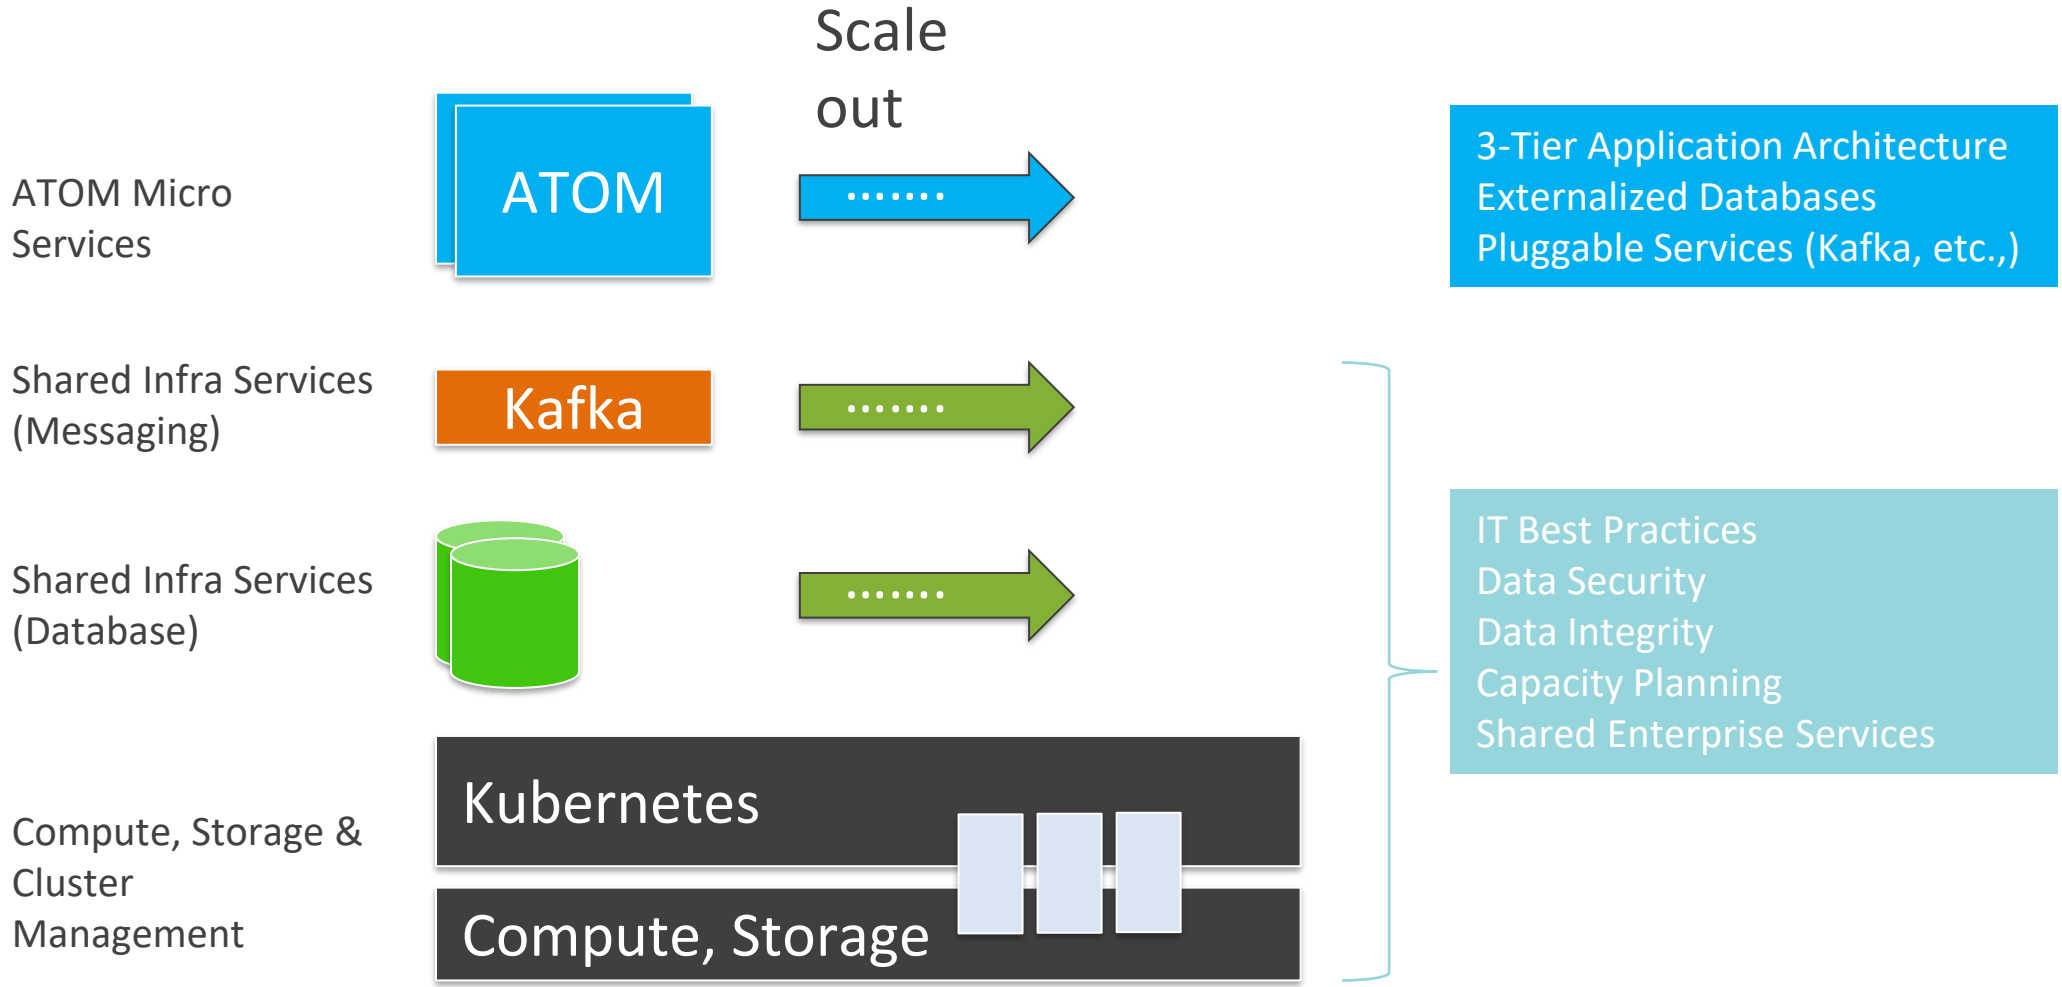


# Anuta ATOM – Software Architecture

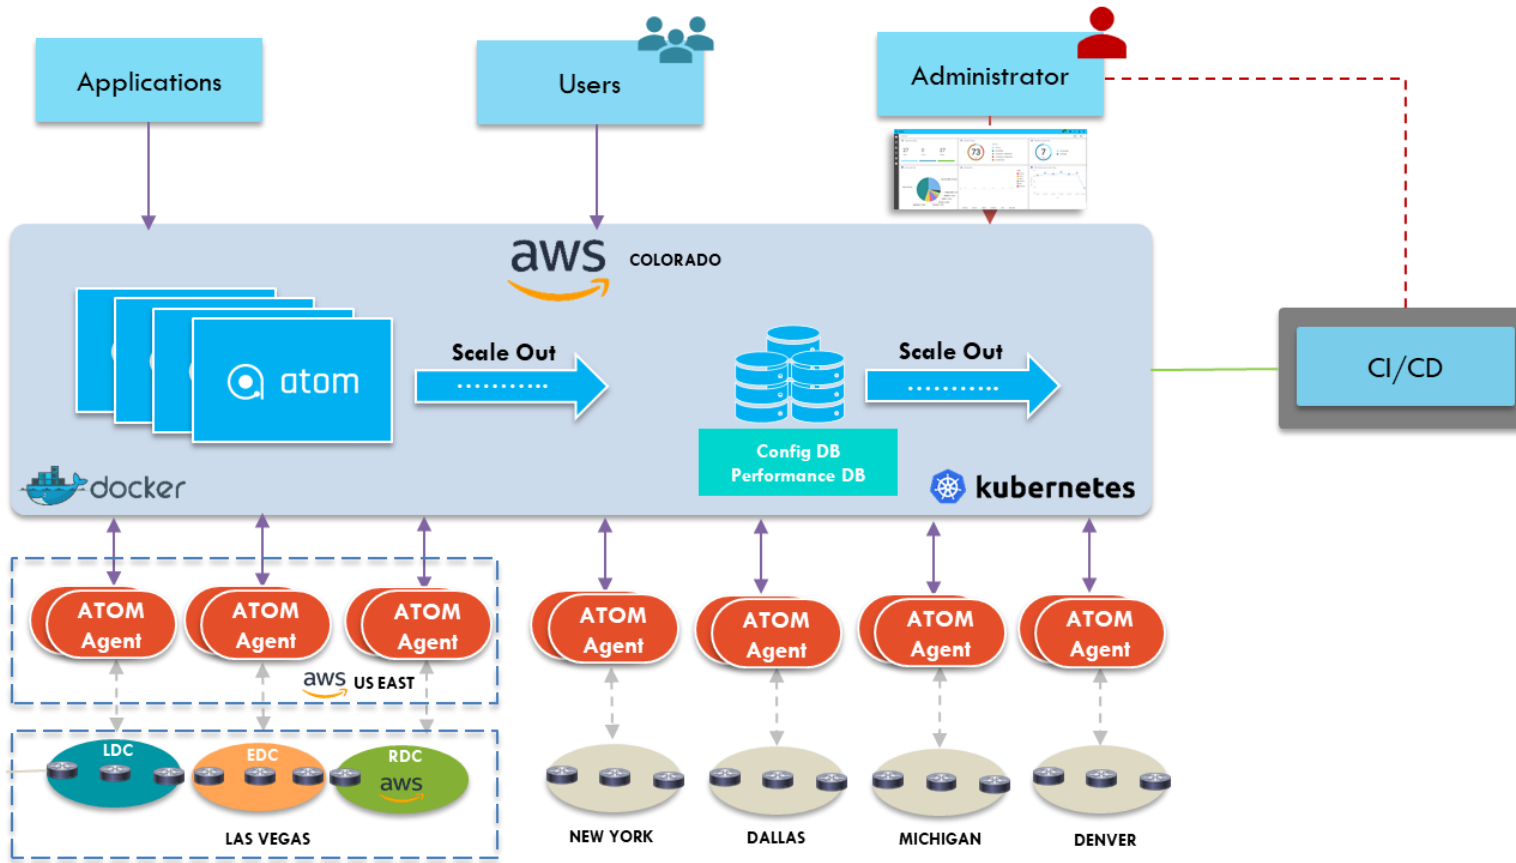


- Scales to Million+ Devices
- Multi-Vendor – 45+ Vendors
- Dockerized & Micro Services Architecture
- Resilient, HA & Multi-Site Distributed Deployment with Disaster Recovery
- Components run Active-Active
- Pluggable Architecture
- SDK Allows Extensions
- KPI Driven Auto-scale
- Opensource Components
- Integrations into AD, LDAP & TACACS
- Model Driven API, UI & Schema browser

# Anuta ATOM - Deployment Overview



# ATOM Deployment in Multi-Availability Zones



- Multi-site distributed deployment
- Multi-Availability Zone deployment in AWS
- Deploy ATOM Agents in Local AWS Zone
- Single-pane-of-glass for all Zones
- Horizontal Scalable architecture



# ATOM Cloud – Agent Life-Cycle



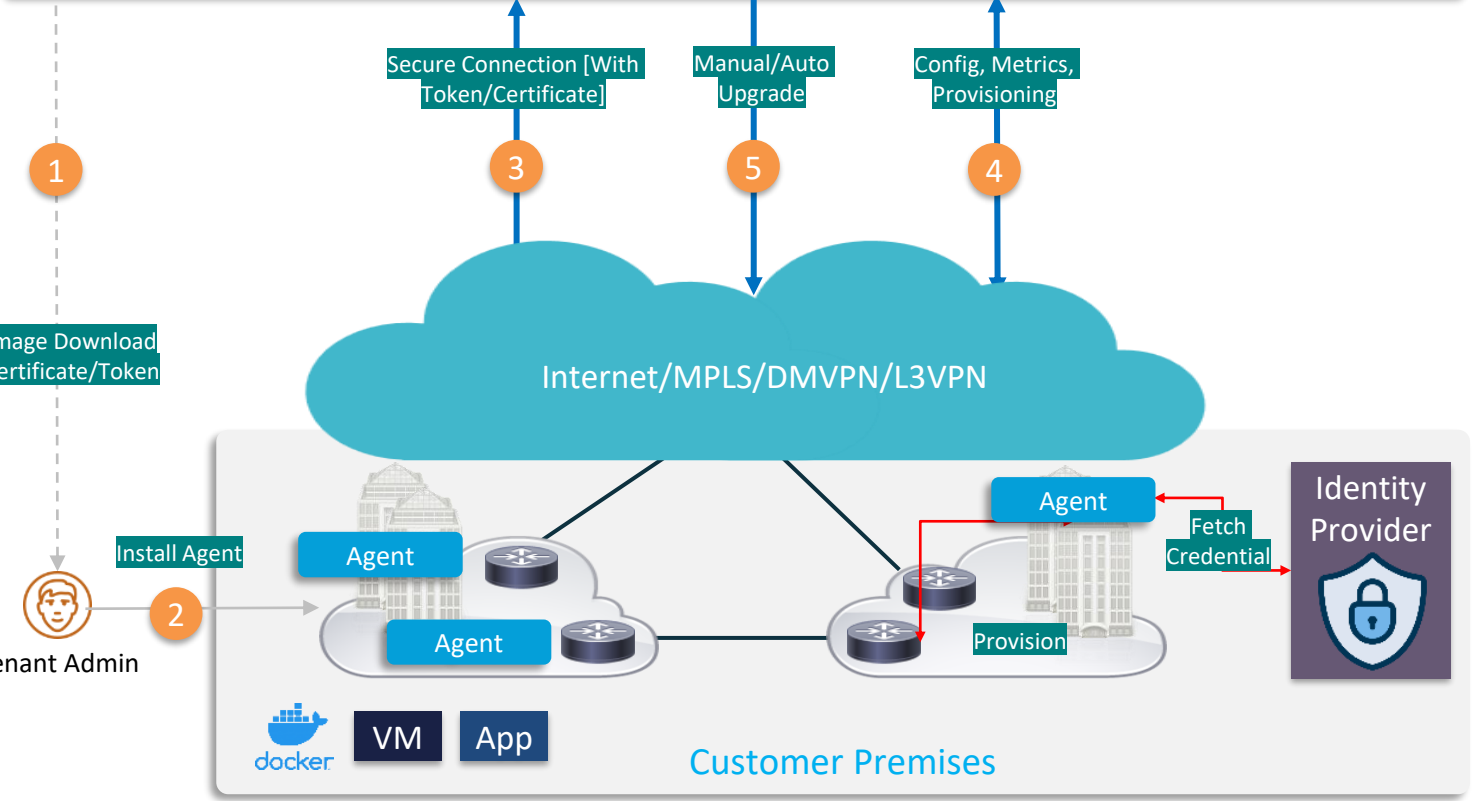
**ATOM Portal**  
 Agent Deployment Files  
 Agent Lifecycle  
 Agent Diagnostics

**ATOM Cloud**

**Agent Manager**  
 (Lifecycle, Connection Mgmt, Certificates etc.,)

Agent    Agent    Agent

**Device Manager**  
 (Connection Mgmt, Configuration Mgmt, Monitoring, Provisioning)



## Hassle-Free Installation

- Self-Service Agent On-boarding
- Multiple Form factors – Docker/VM/App
- Connection Initiation from Agent
- Cloud Deployment or On-Premises

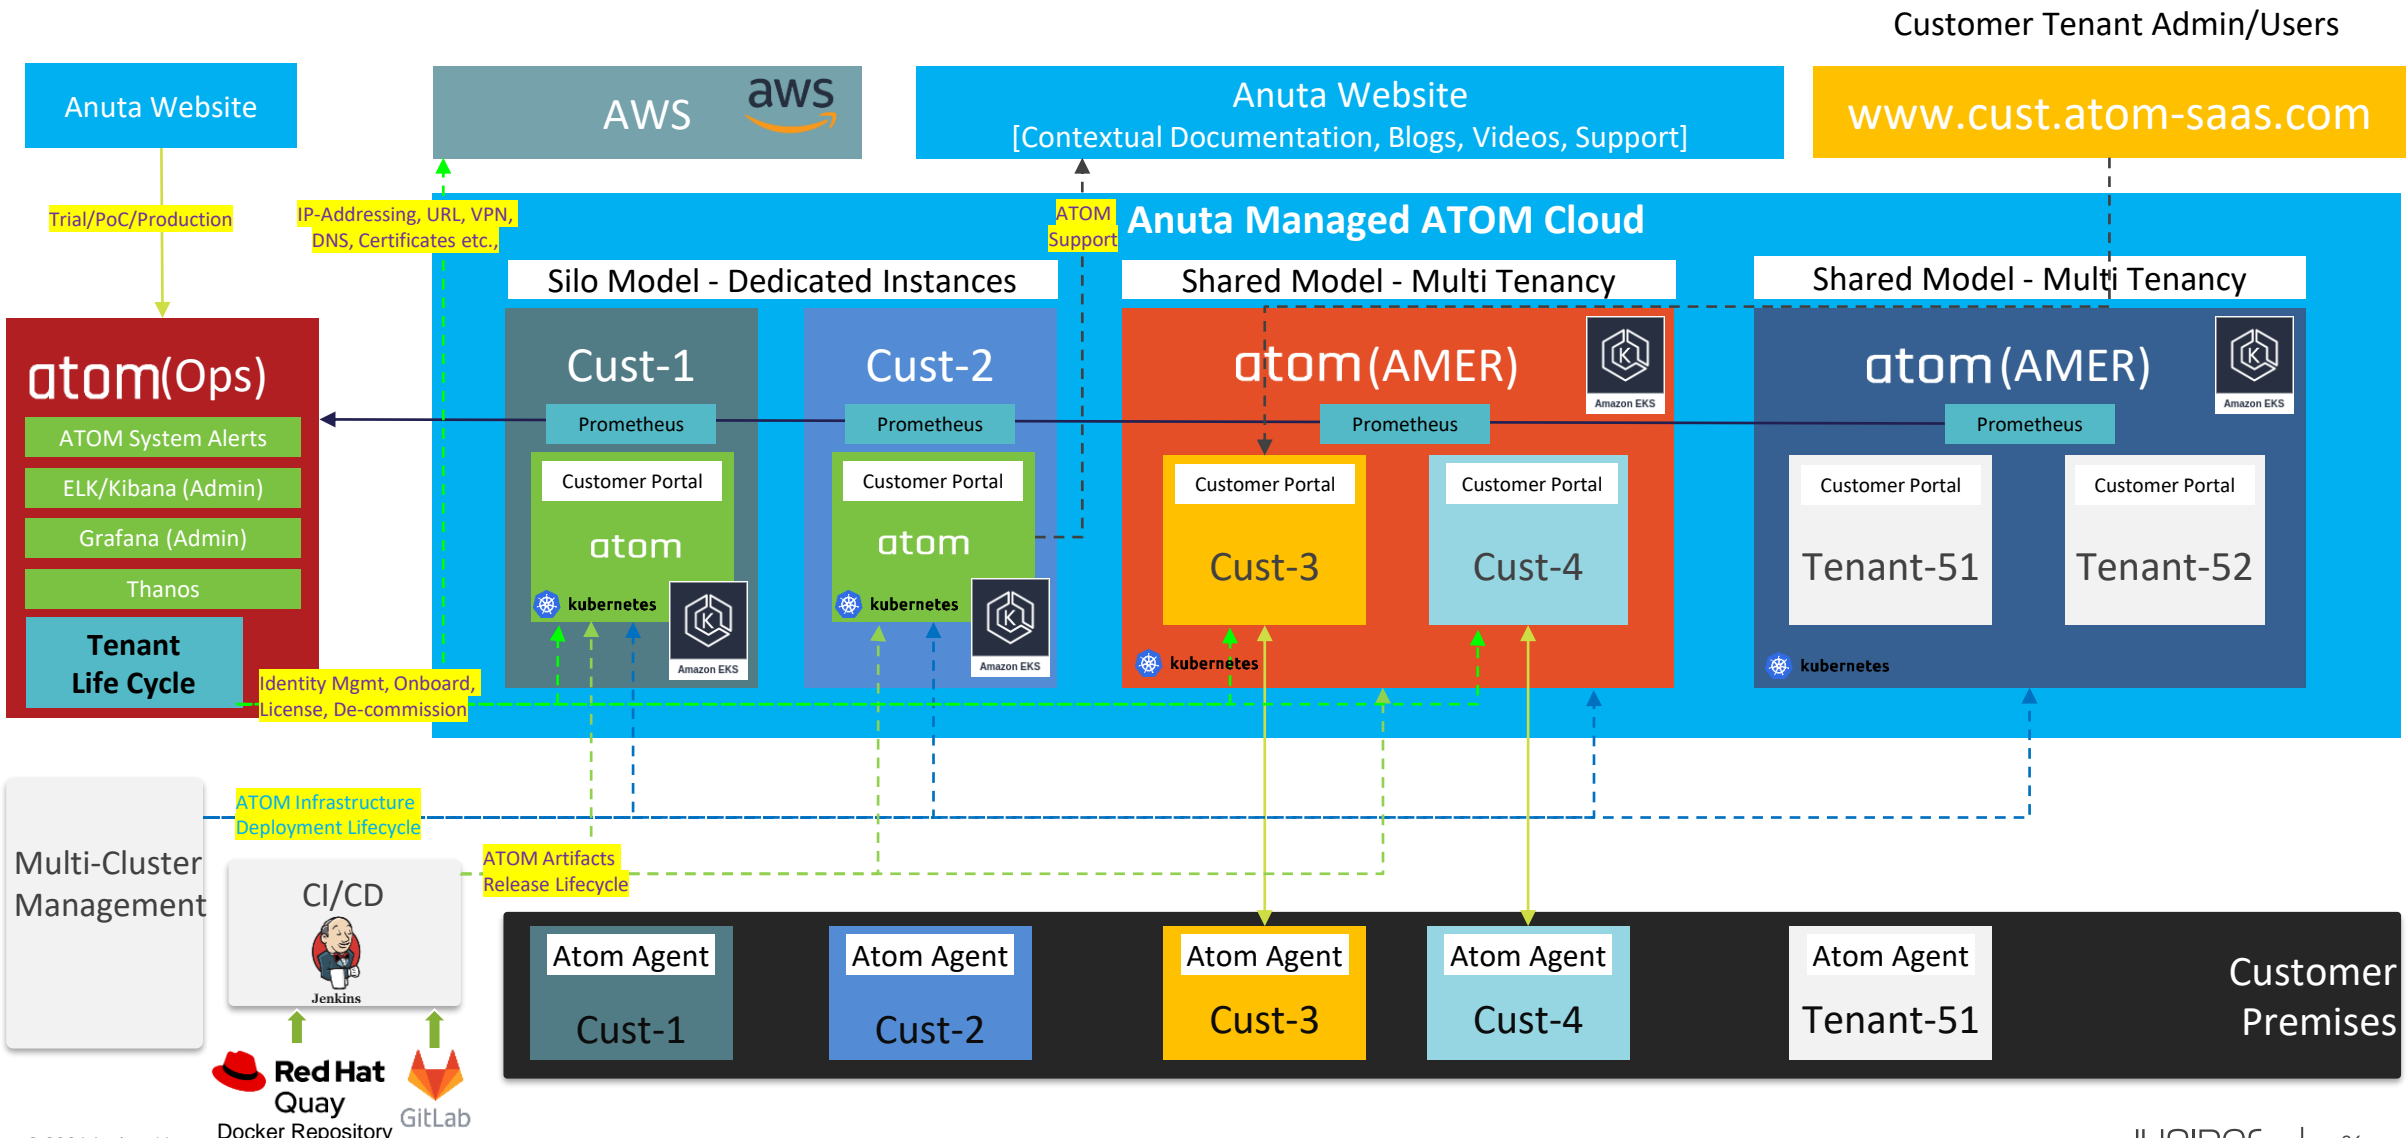
## Security

- Dedicated Per Tenant
- Secure Connection to ATOM Cloud
- Localized Credential Management

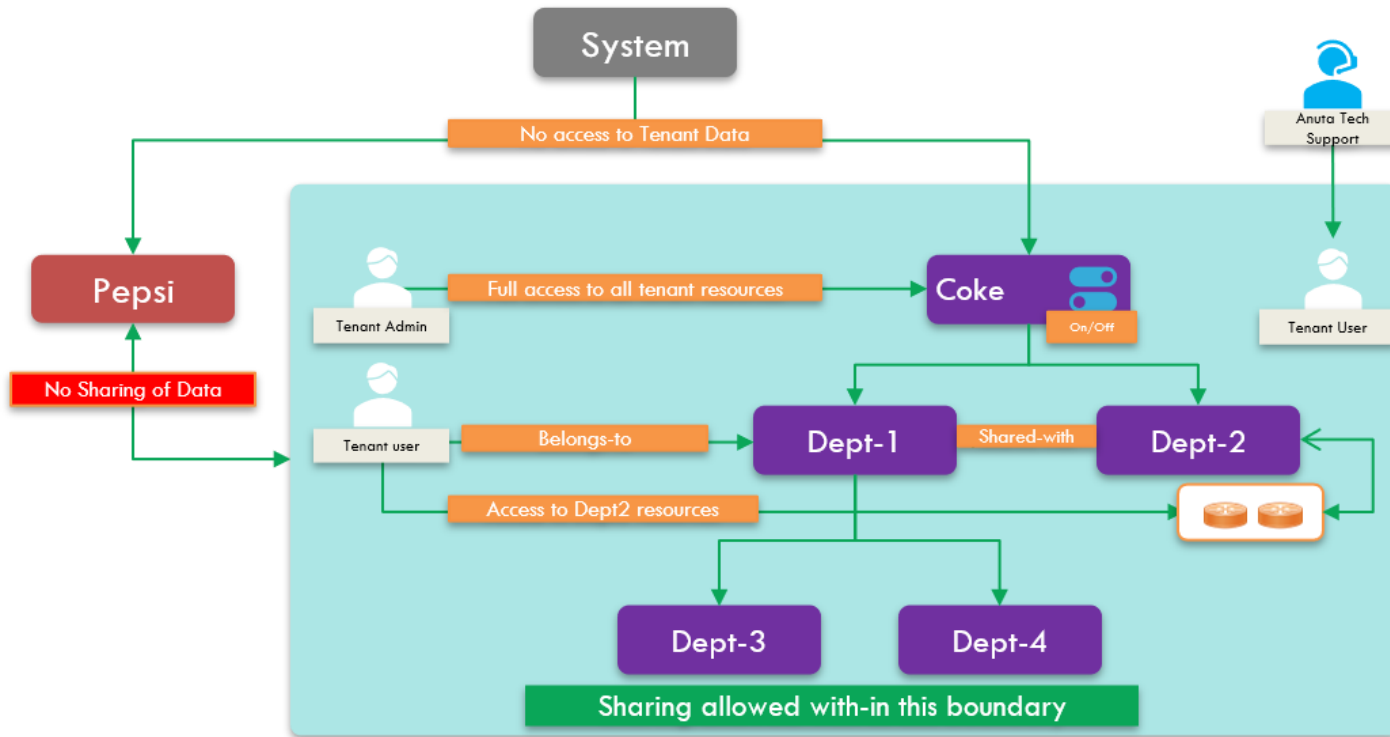
## Easy Maintenance

- Zero-Touch Maintenance
- Stateless & Disposable
- Localized Credential Management

# Anuta ATOM - Cloud Management Reference Architecture



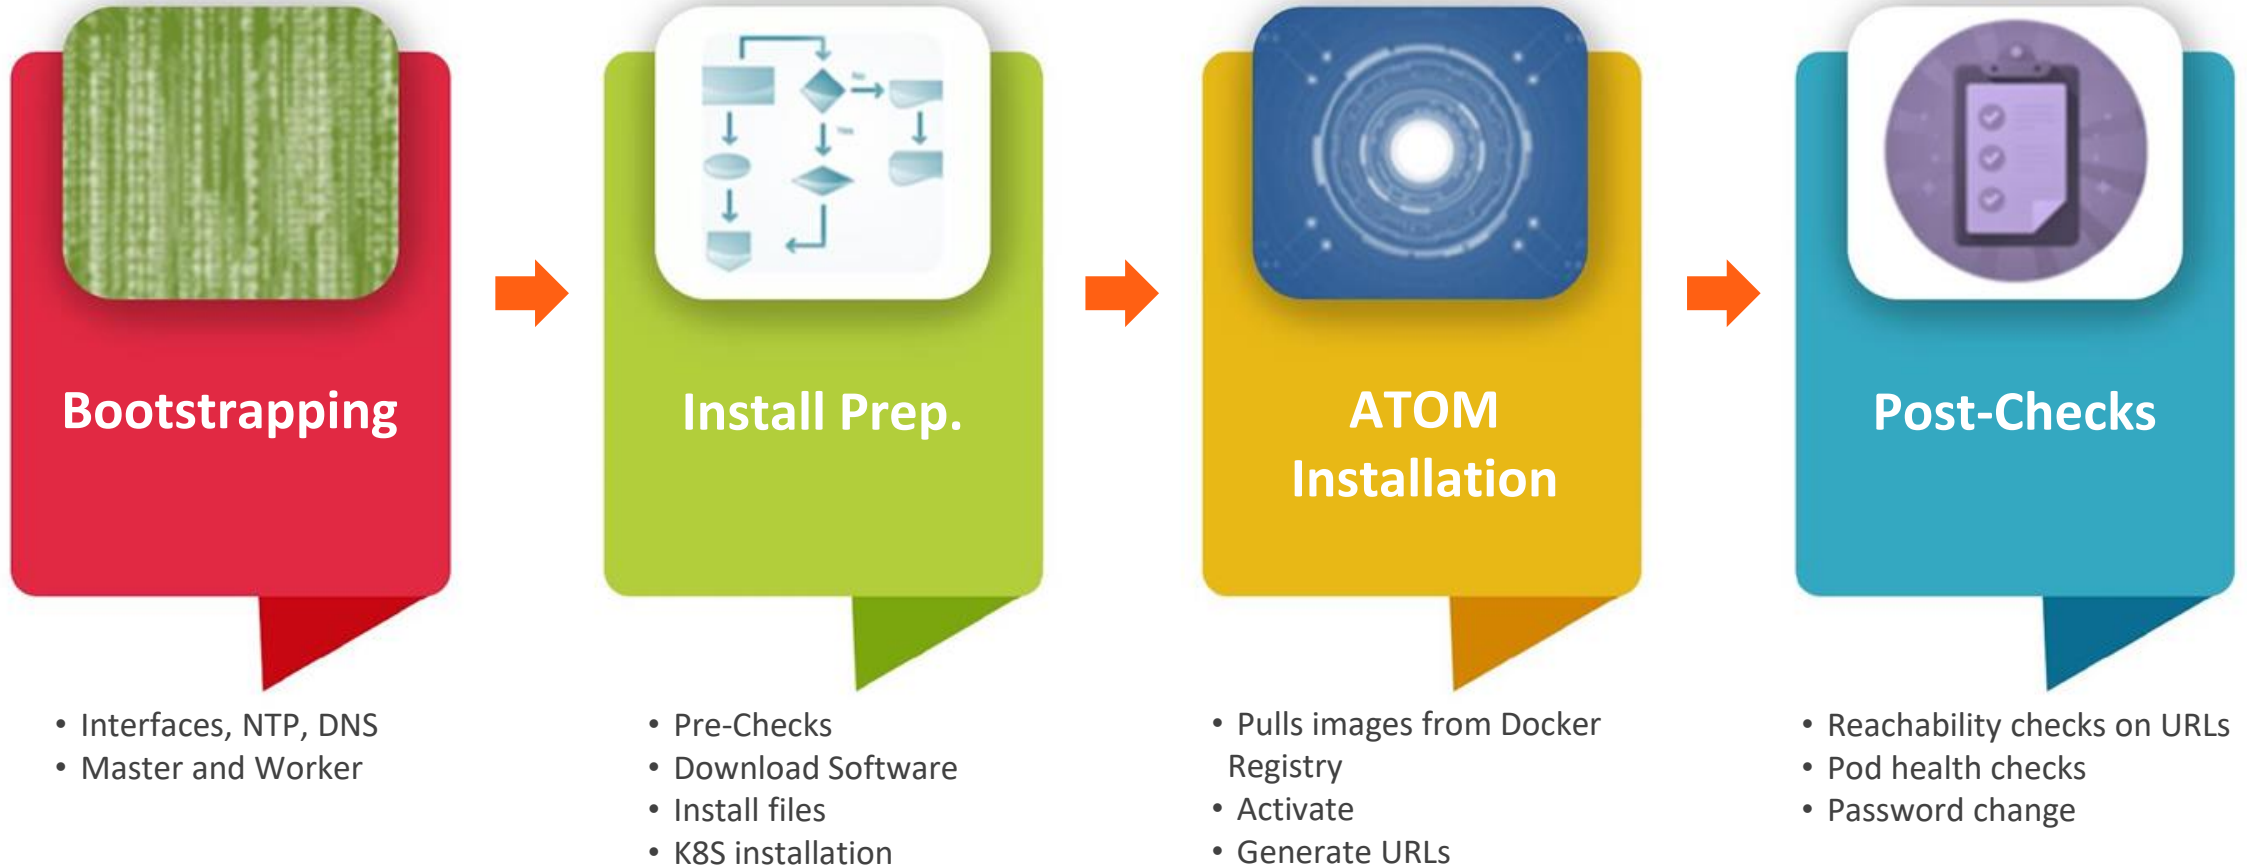
# Tenancy & Organization Model



- No Data Leaking across Tenants
- All Tenant Data Qualified with Tenant-ID
- System does not have any Access to Tenant Data
- Anuta Tech Support can avail Tenant Setup through Temporary Tenant Created and Authorized User Accounts.
- Multi-Level Organization Model
- Resources can be shared between Departments within a Tenant
- Users assigned to Departments
- Recursive Resource sharing can be enable/disabled

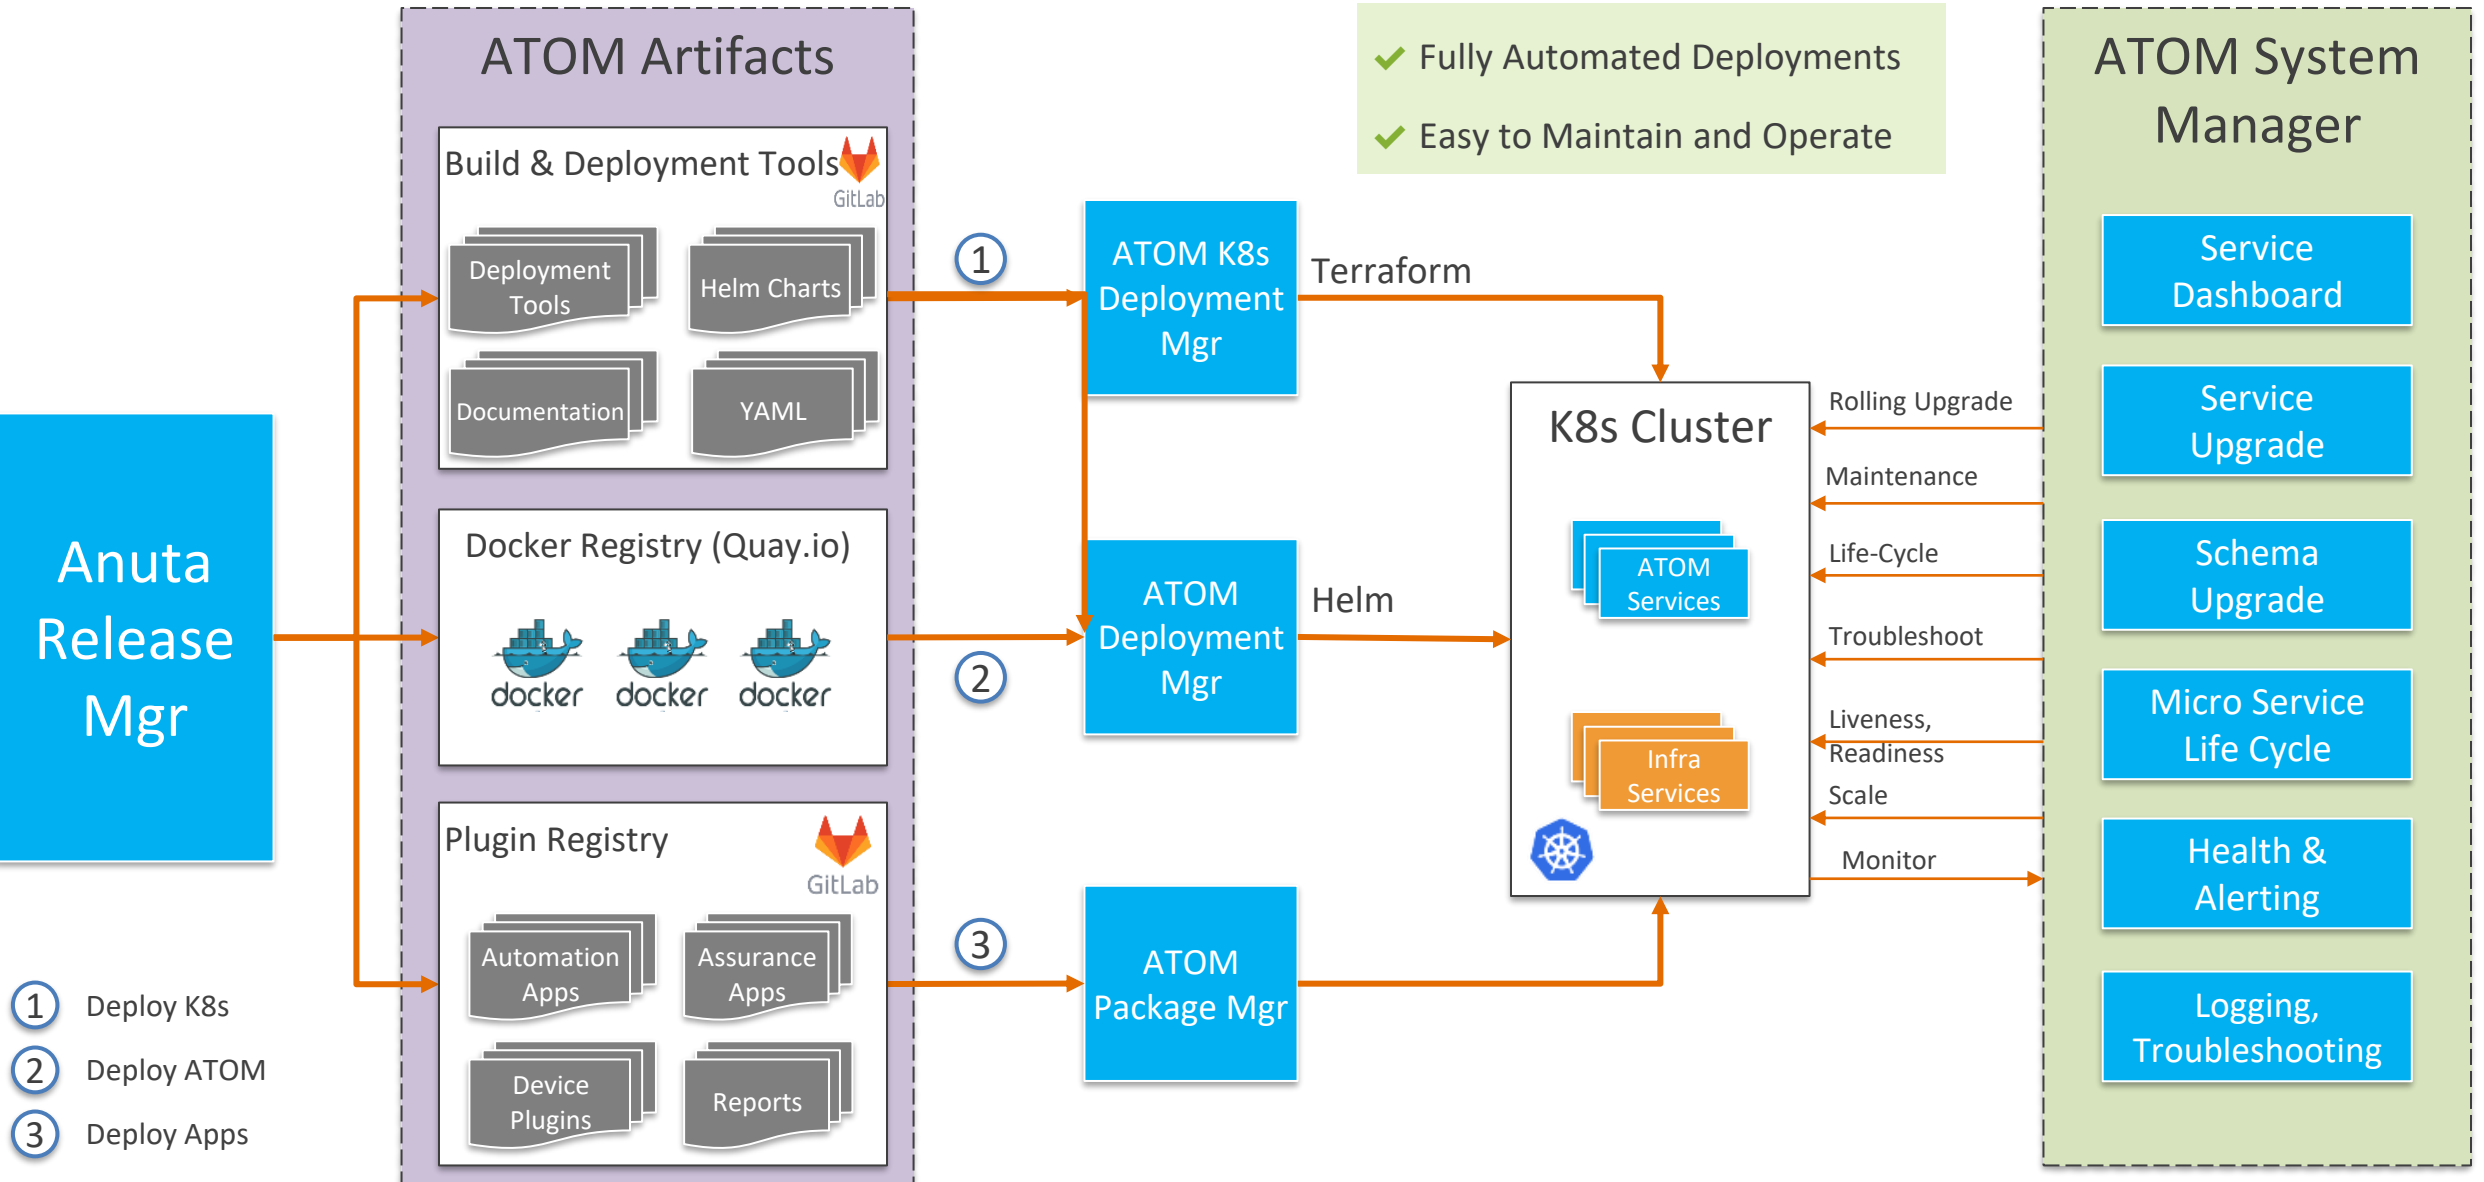


# Anuta ATOM – Automated Installation



Zero-Touch ATOM Deployment within 2 hours.

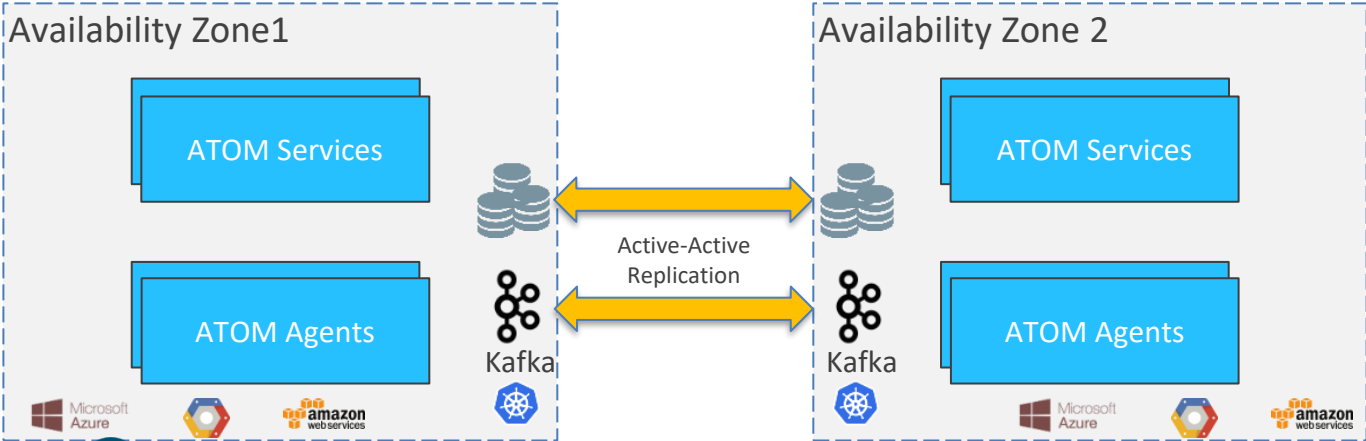
# Anuta ATOM - Deployment, Upgrade & Maintenance



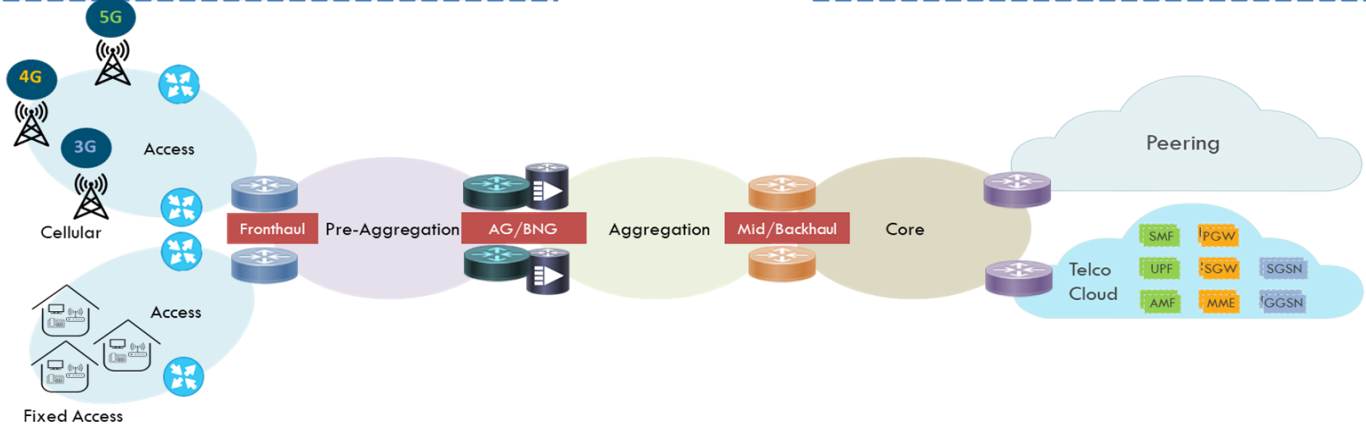
# ATOM Deployment Scenarios – Scalability & HA Architecture



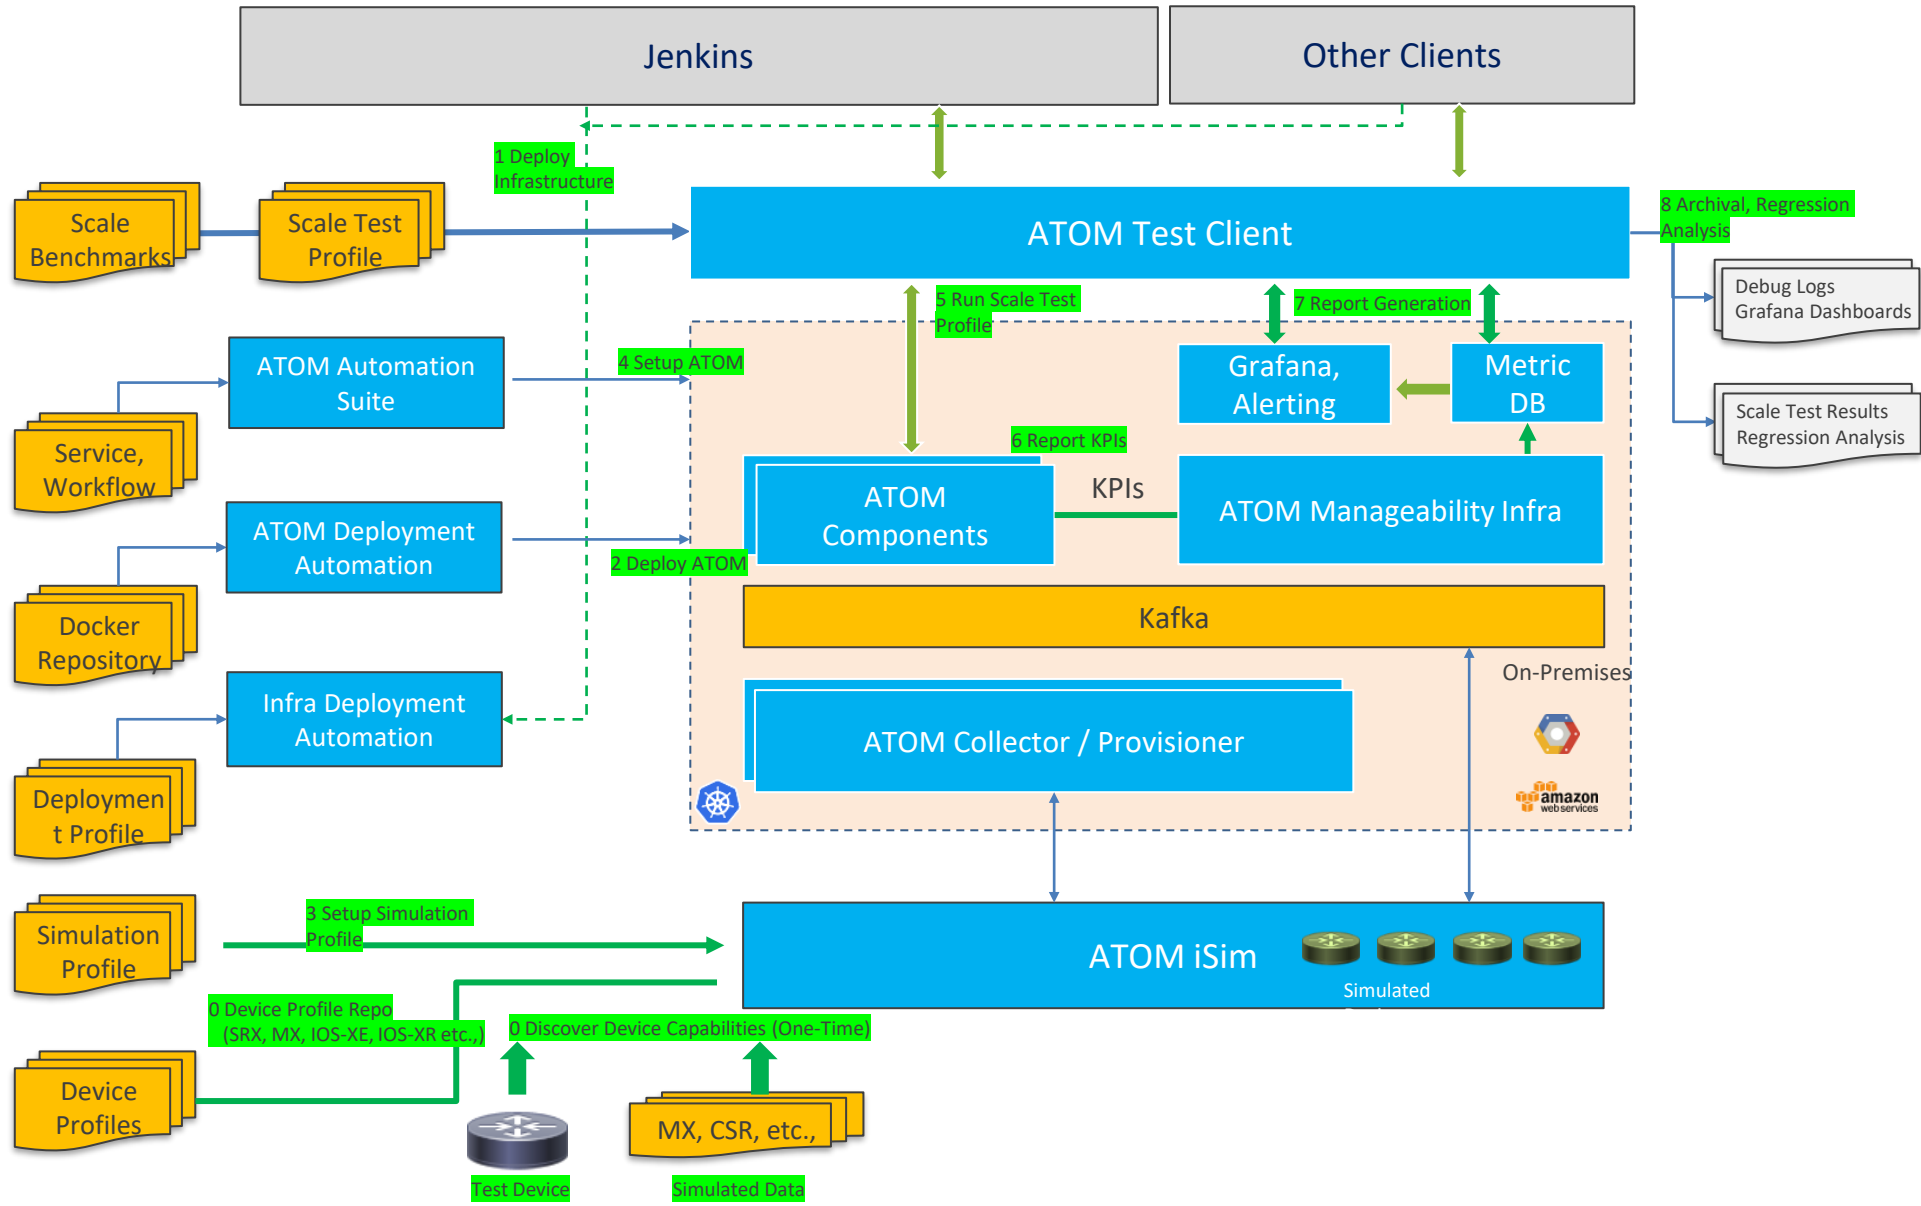
## OSS, BSS, Client



- ✓ On-Premises or CSP Deployments
- ✓ Highly Available / Resilient
- ✓ Multi-Site & Multi-Availability Zone
- ✓ Distributed Deployment
- ✓ Disaster Recovery
- ✓ Active-Active



# ATOM Scale Test Infrastructure



## ATOM iSim

- ✓ Device Configuration – Read, Write
- ✓ Show Command Outputs
- ✓ Exec Commands
- ✓ SNMP Data generation
- ✓ Configurable Response Delays
- ✓ Error Injection
- ✓ Telemetry Data generation
- ✓ Record and Replay
- ✓ Randomize Data
- ✓ NETCONF or CLI

## ATOM Deployment Automation

- ✓ One-Click Automation
- ✓ Compute & Storage
- ✓ On-Premises or Cloud Based

## ATOM Test Client

- ✓ Variation in Load – Controlled, Random etc.,
- ✓ Randomized Inputs
- ✓ Report Generation
- ✓ Regression Analysis

# ATOM Scale - Services

Type	Total Resources	Peak Throughput (Services/Hr)	Total Volume <sup>*3</sup>	Endpoints/subinterfaces/ Devices <sup>*4</sup>	Latency / Service <sup>*5</sup>
S	CPU = 24 vCPUs Memory = 192GB RAM	10-12K	575K	575K/575K/25K	9-10 sec
M	CPU = 56 vCPUs Memory = 450GB RAM	50-60K	3M	3M/3M/140K	9-10 sec
L	CPU = 84 vCPUs Memory = 672 GB RAM	100-120K	5.5M	5.5M/5.5M/300K	9-10 sec
XL	CPU = 160 vCPUs Memory = 1.2TB RAM	180-225K	10.8M	10.8M/10.8M/550K	9-10 sec
<b>Horizontal Scale</b>	As all of these microservices are deployed on K8s, an increase in replica count makes sure the load gets distributed. Scale beyond XL can be achieved by increasing the replica count and observing the throughput, latency. Further details on tuning in the <a href="#">performance fine tuning</a> section.				

Scales to Millions of Services & Thousands of Devices ensuring Throughput & Latency Requirements

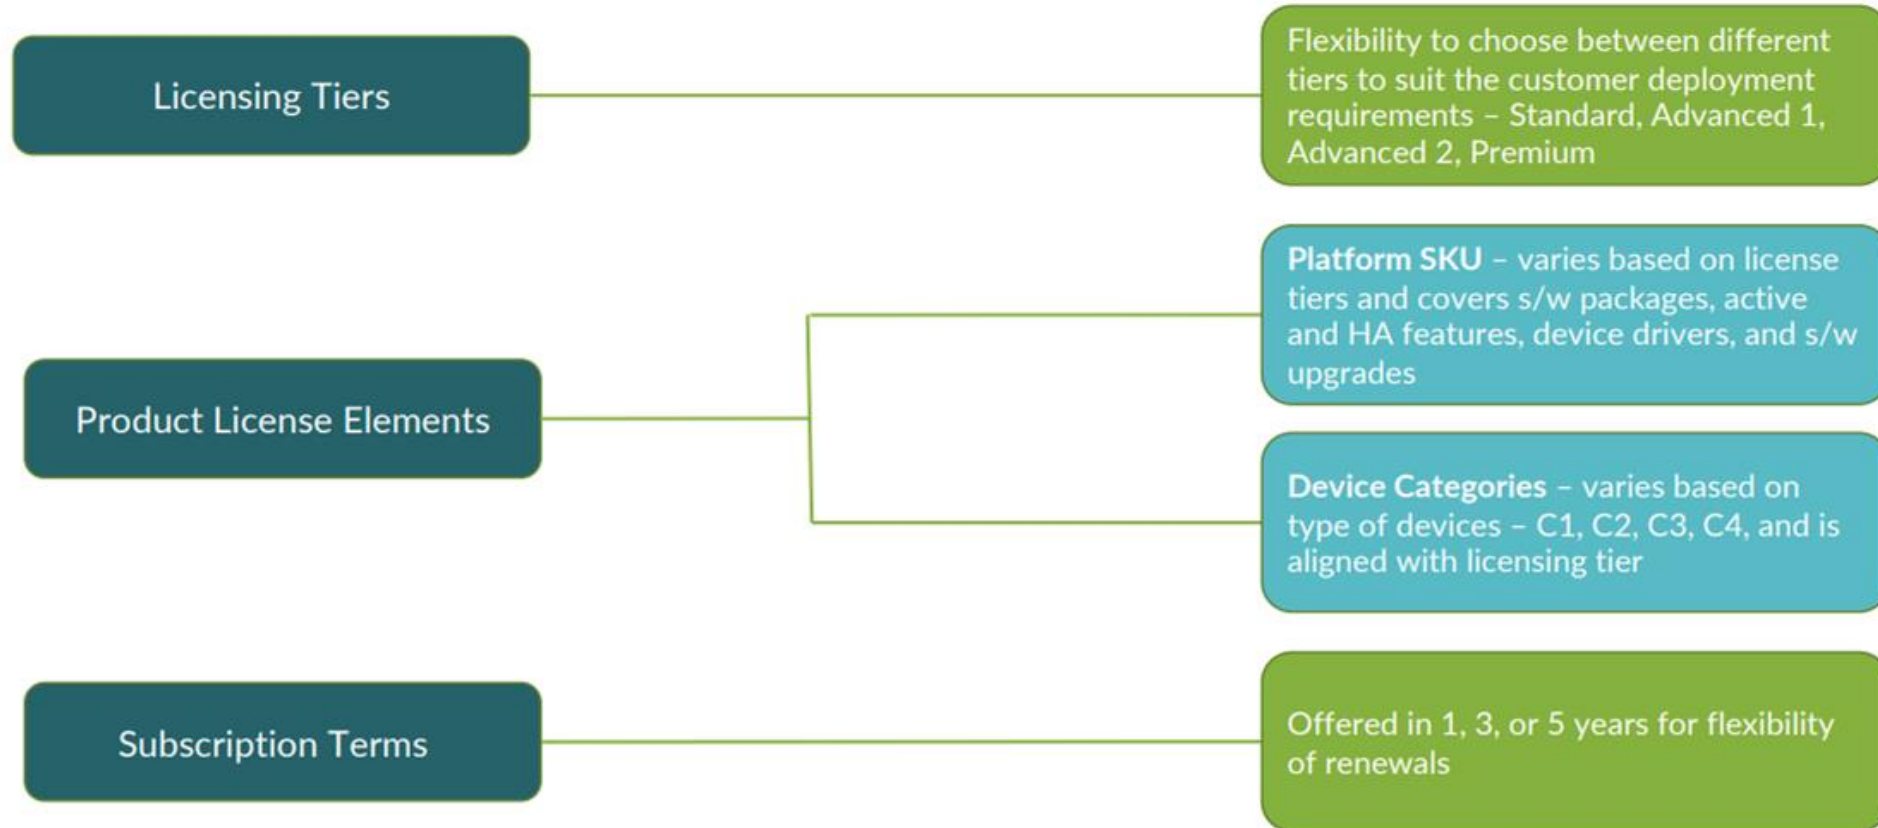


# ATOM Scale - Workflows

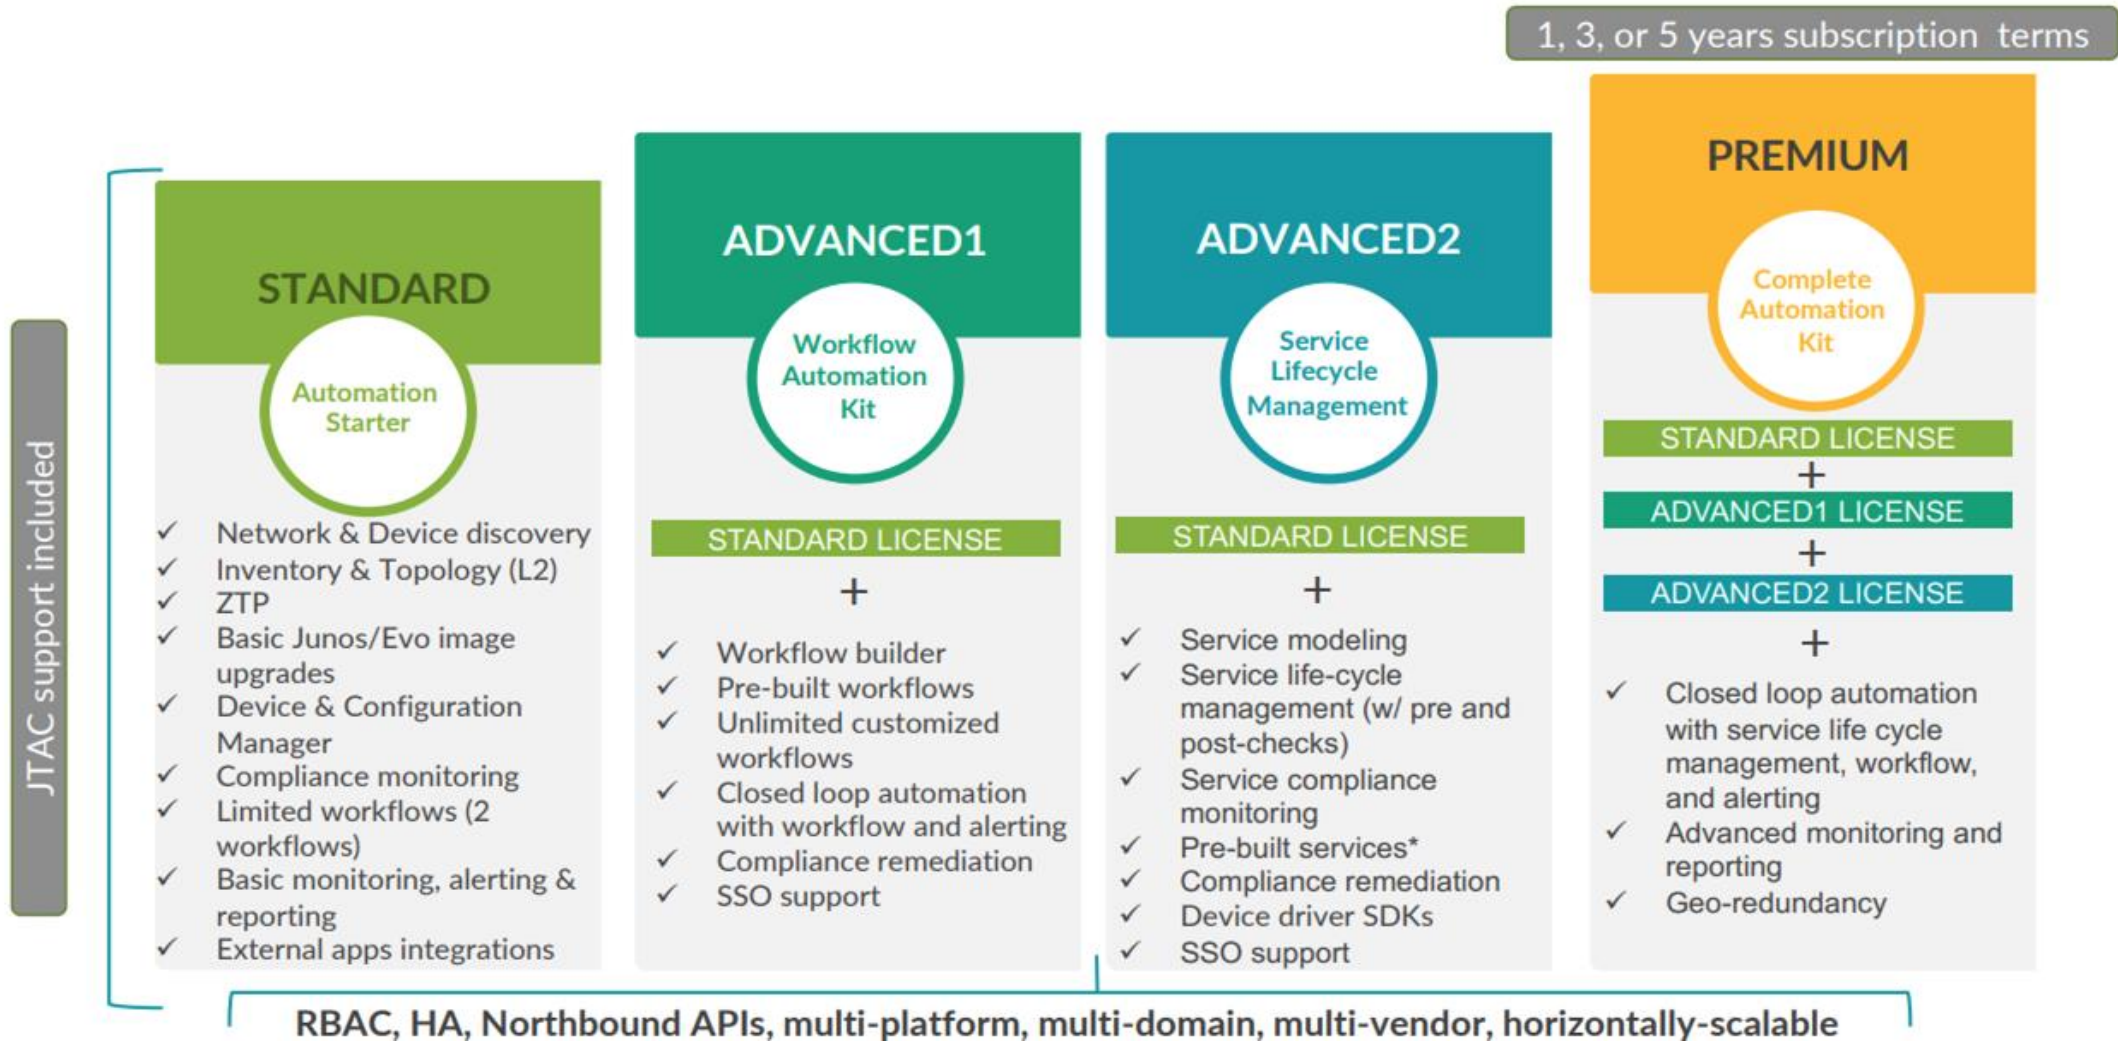
Type	Total Resources	Total Device Count	Time Taken for Device Upgrade	Device Upgrades / Hr	Total Time to Upgrade
<b>S</b>	CPU = 24 vCPUs Memory = 192GB RAM	100	~1hr	40	2hrs 27 mins [100 devices]
		200	~1hr	54	3hrs 40mins [200 devices]
<b>M</b>	CPU = 56 vCPUs Memory = 450GB RAM	500	~1hr	175	2hrs 51 mins [500 devices]
<b>L</b>	CPU = 84 vCPUs Memory = 672 GB RAM	2,500	~1hr	797	3hrs 8 mins [2500 devices]
		5,000	~1hr	898	5hrs 34 mins [5K devices]
<b>XL</b>	CPU = 160 vCPUs Memory = 1.2TB RAM	20,000	~1hr	2962	6 hrs 45 mins [20K devices]

Scales to Thousands of Workflows & Devices ensuring Throughput & Latency Requirements

# Anuta ATOM Ordering Elements



# Anuta ATOM Licensing Tiers




# Device Categories

Category	Description	Examples
C1	extra small: mobile backhaul access type of devices, small/remote CPEs	small ACX, branch small SRX, Ericsson access R6K, ...
C2	small: pizza box switches, TOR switches, metro access switches	fixed form factor EX/ACX/QFX, etc
C3	medium: small chassis switches, small chassis routers, firewalls. Chassis with $\leq 6$ slots	MX10003, PTX10K3, ..
C4	large: large chassis systems. Chassis with $>6$ slots	MX10K, PTX, MX2K, MX960, large SRX, ...


# ATOM Deployment Specs : On-Prem

## Minimal


- Single Kubernetes cluster
- K8s **Master node**\* – 3 x (4vCPU, 32GB RAM and 300GB SSD)
- K8s **Worker node**\* – 1 x (4vCPU, 32GB RAM and 300GB SSD)
- For each remote site - 1 Atom **Agent VM** (4vCPU, 8GB RAM and 50GB SSD)



K8s cluster



Remote site One




Remote site Two


- \***Master Node** - Master nodes will act as worker node also (Shared Master)
- \***Worker Node** – At least 1 Worker Node is needed, one additional Worker Node will help as buffer
- Total IPs : 4 IPs + 3 VIPs = 7 IPs

## Fully Resilient HA


- Resilient Kubernetes Cluster
- K8s **Master node**\* – 3 x (4vCPU, 8GB RAM and 50GB SSD)
- K8s **Worker node**\* – 9 x (4vCPU, 32GB RAM and 300GB SSD)
- For each remote site - 1 Atom **Agent VM** (4vCPU, 8GB RAM and 50GB SSD)



K8s cluster



Remote site One




Remote site Two

- \***Master Node** – These are dedicated Masters
- \***Worker Node** - 9 Worker Nodes gives ideal resiliency in single site also by spreading those across 3 Esxi's/Servers in that site. ATOM can still function with 8 or 7 worker nodes based on number of sites ATOM is deployed on and resiliency aspects needed.
- Total IPs : 12 IPs + 3 VIPs = 15 IPs

# ATOM Deployment Specs: Multi AZ/Sites

## 2 site deployment

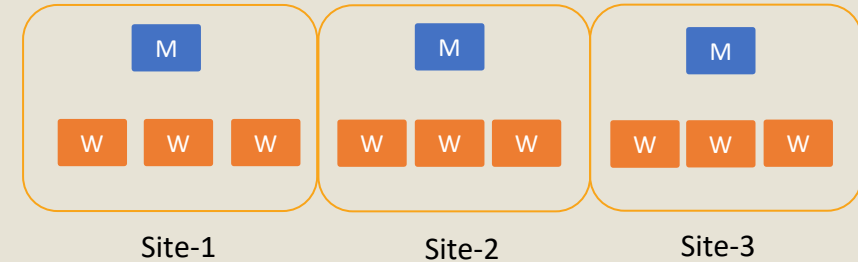
- K8s **Master node** – each site - 2X(4vCPU, 8GB RAM and 50GB SSD)
- K8s **Worker nodes** – each site - 4X(4vCPU, 32GB RAM and 300GB SSD)



- Dedicated Master nodes
- Total IPs : 12 IPs + 3 VIPs = 15 IPs

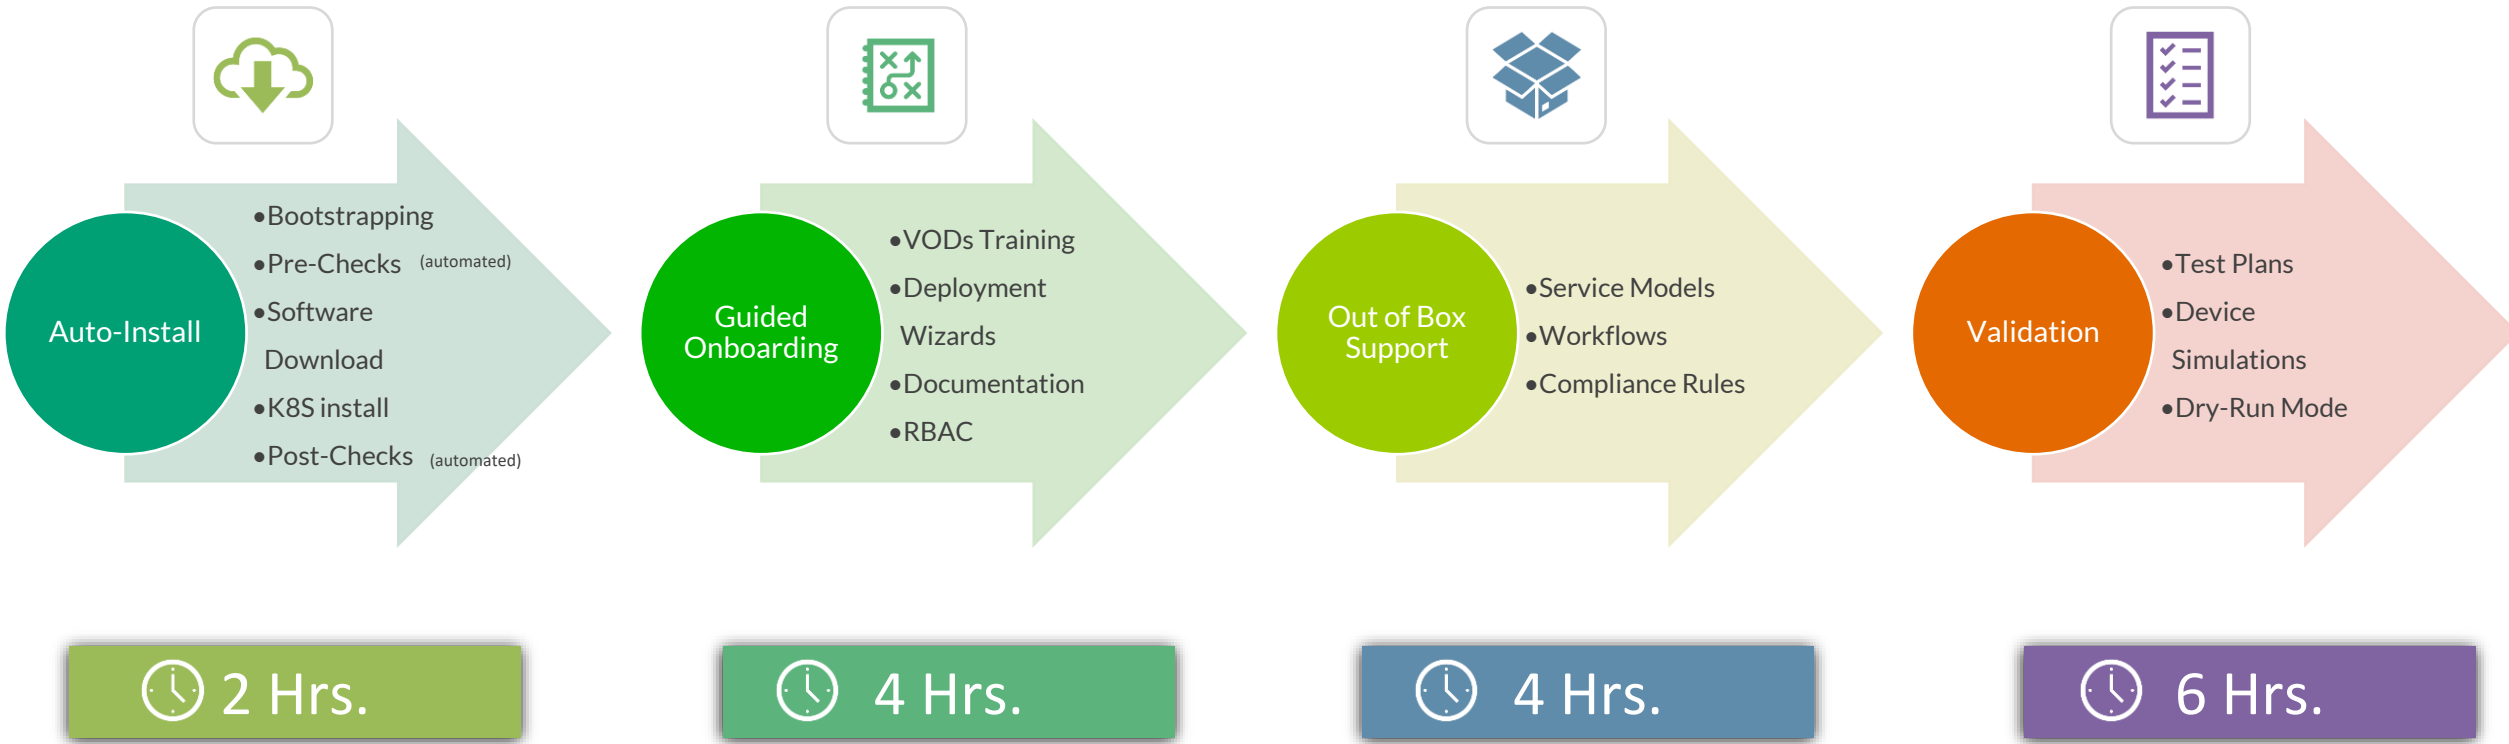
## Fully Resilient 3 site deployment

- Resilient Kubernetes cluster
- K8s **Master node** – each site - 1X(4vCPU, 8GB RAM and 50GB SSD)
- K8s **Worker nodes** – each site 3X(4vCPU, 32GB RAM and 300GB SSD)



- 3 site is the ideal scenario in the cloud
- Dedicated Master Nodes
- Total IPs : 12 IPs + 3 VIPs = 15 IPs

# Anuta ATOM – Time to Demo





# Anuta Trailblazer – Cloud-Native Automation Certifications

## Anuta Trailblazer

### Enhance Your Network Automation skills with Anuta ATOM Certification

Becoming an Anuta Trailblazer helps you ensure you have the technical expertise to efficiently use and maintain the Anuta ATOM platform. Whether you are trying to improve your marketability or impress your customer, Anuta Certification helps you stand out in the crowd.

### Cloud Native Network Automation Associate

#### Introduction to Anuta ATOM Platform

Introduction to Anuta ATOM Platform is a mandatory step in Anuta's training & certificate program. By the end of this course, the individual will be equipped with in-depth theoretical knowledge on Anuta ATOM features and its diverse set of use cases.



The five primary areas covered in this certification course are:

- > Configuration Management Overview
- > Service Orchestration Overview
- > Workflow Automation Overview
- > Compliance Management Overview
- > Monitoring & Closed-loop Automation Overview


[Enroll Now](#)

## Cloud Native Network Automation Professional

### Deep-Dive into Anuta ATOM Operations for Network Operators and Administrators

A deep-dive into operations and administration of Anuta ATOM platform. By the end of this course the individual will have gained a detailed understanding of ATOM architecture, operations and administration.

Prerequisite: Complete the course – Cloud Native Network Automation Associate



The primary areas covered in this certification course are:

- > ATOM Software and Component Architecture
- > Device Discovery & Resource Management
- > ATOM Operations - Service, Workflow, Compliance and CLA
- > ATOM Administration
- > ATOM Debugging & Troubleshooting

[Contact Us for Professional Course](#)

## Cloud Native Network Automation Expert – Workflow

My Profile

9% COMPLETE Last activity on July 7, 2021 12:43 am [IN PROGRESS](#)

### Course Content

Expand All

- ✔ Workflow Overview & Use Cases  
5 Topics | 1 Quiz [Expand](#)
- ✔ ATOM Workflow Constructs  
16 Topics | 1 Quiz [Expand](#)
- ✔ Hands On Exercises  
17 Topics | 1 Quiz [Expand](#)
- ✔ ATOM SDK  
9 Topics | 1 Quiz [Expand](#)
- Final Quiz – Cloud Native Network Automation Expert – Workflow



**Amar Madhavan** • 2nd  
Sr System Engineer at Juniper Networks | Automation & Cloud Architect | 3x JNCIE  
13h •

ATOM Professional!

Check out Juniper's Paragon Automation where ATOM is the Orchestrator!



28 • 4 comments

Like Comment Share Send




Add a comment...

## Cloud Native Network Automation Professional - Compliance Management

### Deep-Dive into Anuta ATOM Compliance Management for Network Operators and Administrators

A deep-dive into defining and executing compliance management in Anuta ATOM platform. By the end of this course the individual will have gained a detailed understanding of ATOM compliance management module.

Prerequisite: Complete the course – Cloud Native Network Automation Professional.



The primary areas covered in this certification course are:

- > Introduction to Configuration Compliance
- > Overview of Configuration Compliance in ATOM
- > Examples of Compliance Policies - CLI and Jinja2 Templates
- > Examples of Compliance Policies - NETCONF/YANG
- > Compliance using ATOM APIs


[Contact Us for Professional - Compliance Course](#)

## Cloud Native Network Automation Expert - Workflow

### Hands-on Experience with ATOM Workflow

A deep-dive into the Workflow Automation module in Anuta ATOM. By the end of this course, the individual will understand all the powerful capabilities to automate the complex methods of procedures.

Prerequisite: Complete the course – Cloud Native Network Automation Professional



The primary areas covered in this certification course are:

- > ATOM Workflow Overview and Use-Cases
- > Workflow Constructs - Tasks, Events, Forms
- > 15 Hands-On Use-Cases
- > ATOM SDK to build Service Packages
- > Integrating Workflow with a Service Package

[Contact Us for Expert - Workflow Course](#)

Self-Paced Training

100+ VODs

Enroll at:

<https://www.anutanetworks.com/anuta-trailblazer/>



The background is a night cityscape with a green overlay. A network diagram is overlaid on the city, consisting of numerous vertical lines with glowing green dots at the top, and a complex web of white lines at the bottom representing a network or data flow. The text "Thank you" is centered in the green area.

# Thank you

---

JUNIPER NETWORKS | Engineering  
Simplicity