Ventech Solutions

Ventech Project Showcase:

AIOps Dashboard (The Dashboard) and Ventech AI/ML Capabilities

May, 2023

Table of Contents

List of Figures

List of Tables

ACRONYMS

ADO	Application Development Organizations
AI	Artificial Intelligence
AIOPS	Artificial Intelligence in IT Operations
API	Application Programming Interface
ARIMA	Autoregressive Integrated Moving Average
BDD	Behavioral Driven Development
BERT	Bidirectional Encoder Representations from Transformers
BOW	Bag of Words
CBOW	Continuous Bag of Words
CI/CD	Continuous Integration/Continuous Deployment
DRL	Deep Reinforcement Learning
LDA	Latent Dirichlet Allocation
LOB	Line of Business
LSTM	Long Short-Term Memory
MAPE	Mean Average Prediction Error
ML	Machine Learning
MLOPS	Machine Learning Operations
MVP	Minimum Viable Product
NLP	Natural Language Processing
POC	Proof of Concept
RBAC	Role-Based Access Control
ROI	Return on Investment
SaaS	Software-as-a-Service
SDLC	Software Development Lifecycle
SHAP	SHapley Additive exPlanations
SOS	Service Operations Support
TCN	Temporal Convolutional Network
TDD	Test Driven Deployment
VAR	Vector Autoregressive Models

PART 1:
INTRODUCTION

The AIOps Operations Dashboard ("The Dashboard") is a product we developed at the request of a Government Agency (AGCY) for an enterprise monitoring dashboard. It provides a layer of data visualization into the health and system status (up, down, or degraded) to 66 often disparate IT services within their environment, supporting the application systems infrastructure and data center support application, Application Development Organizations (ADOs) and various Lines of Business (LOBs). The Dashboard provides that layer of data aggregation in a single pane view, incorporating data analytics.

There are a total of 66 services from the environment into The Dashboard, eight of which have been enabled with Artificial Intelligence (AI)/Machine Learning (ML) capabilities for up to 30 minutes of uptime prediction and anomaly detection, with several more capabilities planned for 2023 and beyond.

1. BENEFITS, VALUE PROPOSITION, AND PRODUCT ROADMAP

Our joint vision with the Government Agency Department (DEPT) is to have The Dashboard deliver and maintain a single pane of glass for all IT services while also providing the following benefits:

- Promote **transparency** and **collaboration**.
- Build environment **trust**.
- Provide service level **monitoring**.
- Translate technical topics (raw metrics and logs) into understandable language (non-technical meaning/comprehension).
- Enable cultural change.
    - Transition to a **proactive** culture rather than a reactive one.
    - Enhance accountability and encourage transparency.

Our **unified dashboard** provides value to AGCY executives and service owners by:

- Supporting real-time system health status.
- Providing ML capabilities such as:
    - **Predictive analytics:** Advance warning of potential service issues up to 30 minutes into the future and providing the ability to take preventative action.
    - **Anomaly Detection:** Detecting abnormalities within historical data and detecting outliers that point to abnormal system behavior that could point to future recurrence and minimize false positives.
- Reducing environment system downtime.
- Alerting service owners to act proactively.

The Dashboard achieves these features via existing functionality, which includes the Service Availability dashboard, Service Availability reporting, Service Detail drill-down dashboards, Service Level alerts, and a DevSecOps metrics Application Programming Interface (API).

Planned Future Capabilities: We plan to integrate a cost and security information dashboard and Software-as-a-Service (SaaS) service monitoring.

Ventech Solutions' overall roadmap vision has AI/ML expansion in the DEPT's environment as an extension of The Dashboard, as illustrated in **Figure 1**.

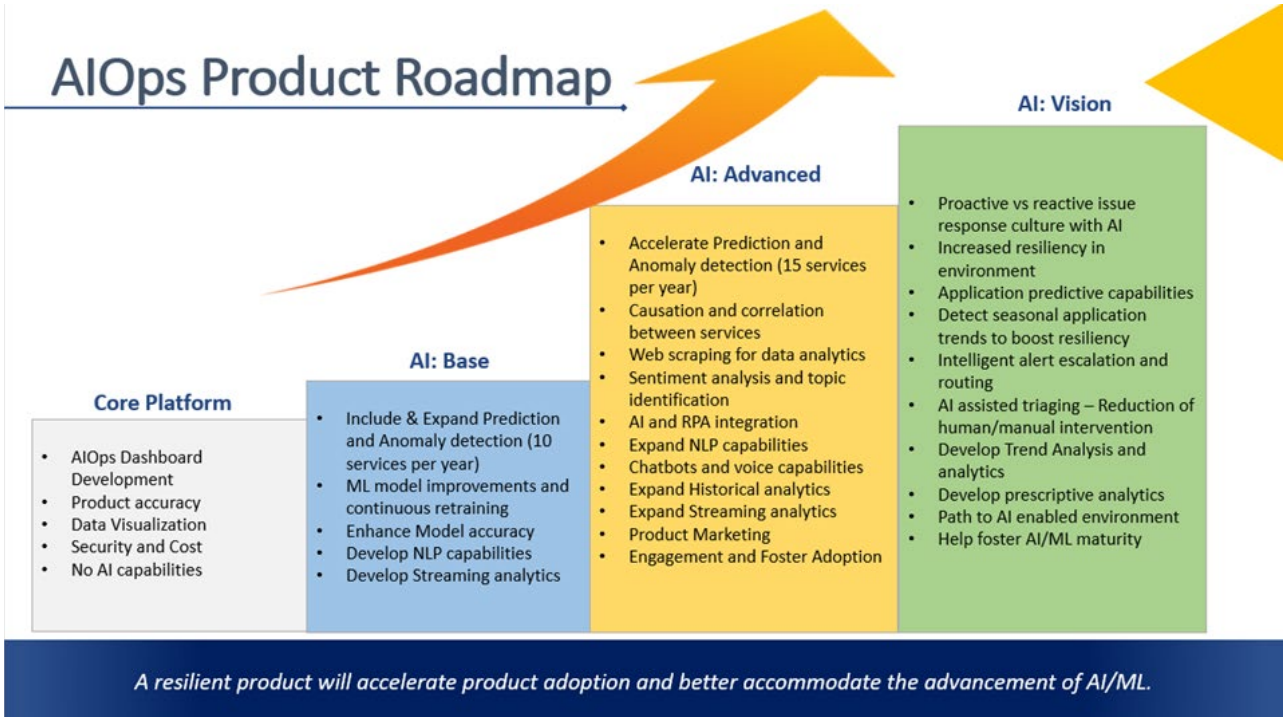

Figure 1: The Dashboard Roadmap and AI/ML Vision

2. DELIVERY STRUCTURE

The Dashboard product is the outcome of a combination of software engineering, development, and AI/ML capabilities. Ventech Solutions formed a software development lifecycle (SDLC)-based organization to develop The Dashboard with carefully selected software engineers to help mature the product. Our purpose-built SDLC-based software and AI/ML engineering team can potentially be used on other LOB applications.

The Dashboard and AI/ML delivery at Ventech Solutions has been organized into three scrum teams, each with a unique focus, the combination of which provides delivery capabilities provided in this document.

1. A core software development and engineering team continues to mature The Dashboard into what will become The Dashboard 2.0, a fully redeveloped product from Version 1.0—the product in operation today. A Minimum Viable Product (MVP) release of Version 2.0 is in live use as of May 2023.
2. An AI/ML scrum team whose purpose is two-fold:
    a. Enable a core set of AI services on The Dashboard to maintain the value of the currently active AI/ML capabilities and expand to other services on The Dashboard.
    b. Drive AI enablement with a larger set of objectives (Figure 1 product roadmap) toward the AI/ML vision that Ventech Solutions proposes.
3. A production support team, whose focus is maintaining the current production version of The Dashboard 1.0, that provides production support and onboarding of new services onto the existing platform today.

The Dashboard and Artificial Intelligence in IT Operations (AIOPS) team is currently highly focused on The Dashboard 2.0 and AI/ML capability objectives; this is a dedicated team on a mission, having achieved significant results and outcomes for AGCY and will continue to do so into the future.

3. CORE PRODUCT CAPABILITIES

This section outlines the work we are currently performing to maintain business continuity of The Dashboard and the operations work to maintain the product.

3.1 The Dashboard: Core Product – Operations, Maintenance, and Expansion Capabilities

At its core, The Dashboard includes the following operations, maintenance, and expansion capabilities:

1. **Sustainability**
2. **Internal Product Support**
3. **Product Engagement and Adoption**
4. **Integration/Expansion**

Operations, Maintenance, and Expansion Overview:

1. Sustainability: Security and Compliance

- **Patching** – The Dashboard resides in the Self-Service Cloud. The Operations and Support team shall be responsible for:
    - Applying AMI updates to The Dashboard infrastructure, such as EC2 instances every 14 days as part of regular deployment schedule.
    - Applying AMI updates to The Dashboard development infrastructure, such as the EC2 instances hosting SonarQube and FluentD at least once every 14 days.
    - Applying patches and updates to development tools, such as SonarQube and FluentD, as necessary, to remediate vulnerabilities.
    - Updating operating systems (RHEL 7) on Jenkins agent EC2 instances via package manager at least once every 14 days.
    - Applying patches or updates to Jenkins agent and other components, as necessary, to remediate vulnerabilities.
    - Remediating Tenable compliance findings (e.g., file system permissions, firewall policies and networking configuration) using the steps and documentation provided by Tenable (varying widely from case to case).

- Remediating Amazon Web Services (AWS) security and best-practice findings in Security Hub and Trusted Advisor using instructions and documentation provided by AWS (varying from case to case).

2. **Internal Product Support:** The Dashboard provides the following internal support activities:

- **Upgrade and Maintenance** – We release new versions of The Dashboard from sprints to ensure that The Dashboard users have continued access to the latest product features and enhancements, as well as issue fixes, including:
  - Supporting The Dashboard 1.0 day-to-day maintenance, continued operation of 66 services and health indicators and maintaining business continuity.
  - Updating new version releases with each sprint to add additional key performance indicators (KPIs), removing subsystems or KPIs to existing services, onboarding new services and resolving issues.
- **The Dashboard Service Issues:**
  - We provide support for any outages and The Dashboard service interruptions.
  - Adhere to AIOPS and Service Operations Support (SOS) Runbook and The Dashboard support.
  - We have developed a 'close the loop' workflow/process to identify and ensure the sequence of events from a health status indicator is followed through to completion to ensure health status on The Dashboard indicates a given system is up and running.

3. **Product Adoption**

- **ADO Engagement:** We engage ADOs in a cooperative community to help The Dashboard provide a "single `pane of glass" view across various environment services in the environment.
- **Future State:** We will build a The Dashboard engagement and operating model to establish and increase engagement and adoption, partnering with other contracted **organizations** such as program management and collaboration services to increase awareness, value and buy-in from ADOs on the usefulness, importance and consumption of The Dashboard.
- **The Dashboard Engagement Model will cover the following topics:**

  - Determine engagement model for ADO.
  - Define onboarding and training for ADO service owners.
  - Build necessary work plans for ADO service and support.
  - Determine ADO interaction.
  - Measure The Dashboard consumption and adoption.
- **Socialization/Marketing:** Share developed knowledge base, processes and tools to support and encourage the use of predictive and anomaly detection capabilities and The Dashboard across the ADOs.

4. **Integration/Expansion: We will integrate with additional ADO applications and monitor key components. We reconcile recurring baseline to ensure a single version of truth.**

3.2 The Dashboard: Core Product Maturity

As a part of continuous product development and improvement, to support future state, a "**The Dashboard 2.0 Core** Platform Development effort" is currently well underway with the MVP in production use. This improvement is building upon the concepts created in The Dashboard 1.0 with an overhaul, adapting the lessons learned from The Dashboard 1.0 to provide a broader platform for future growth and scalability.

We are undertaking a **design-centric persona-driven Agile product management approach** with whole product **design thinking**. The Dashboard 2.0 is implementing modern N-tier architectures, decoupled microservices-based design solutions, enhanced code management and reusable Software Design and DevOps patterns. Core objectives for this effort include an improved product built with the software engineering team already in place and improving upon lessons learned from The Dashboard 1.0— a forward-thinking system scalable to AGCY executive leadership and product owners, ADO DevSecOps, ADO developers, helpdesk and end-users.

The Dashboard 2.0 fundamental themes include **Scalability, Extensibility, SDLC Framework, Usability, Accuracy, and Reliability**. We achieve these themes with several architectural improvements, such as a move to an N-tier microservices-based architecture, a software development lifecycle (SDLC) framework, the establishment of a core software engineering and development team and a software support team for production support of The Dashboard 1.0.

Usability: To improve usability and ADO adoption, we propose a user experience redesign. Rather than continuing use of the out-of-the box open-source dashboard product (Grafana), we suggest a more flexible and extensible solution. Major usability concepts to consider are:

- Testing design concepts with end-users before development begins.
- Explaining metrics displayed and how they are calculated.
- Sharing essential information quickly so users can find it.
- Building with the AGCY design system to create a cohesive environment.

Accessibility: Easy accessibility for users with disabilities (Section 508 Compliance):

- Starting with accessibility guidelines instead of building them later.
- Ensuring enough contrast between elements for users with visual impairments.
- Designing visualizations that are understandable without color to enable usability for colorblind users.
- Using alternative text, skip navigation, reading order and (Accessible Rich Internet Applications) ARIA content so information is available to users navigating with screen readers.
- Including captions for any video content.

Relevance: We continuously follow Agile Product Management best practice approaches. We adopt a persona-driven, user empathy-based approach to influence product design to ensure it is relevant to what primary and secondary personas desire in a product.

We follow these best practices:

- Continue engagement with existing and potential users to understand their needs.
- Conduct persona interviews across multiple product lines to influence product design.
- Develop features based on feedback from interview participants.
- Use Human Centered Design (HCD) practices and empathy maps to guide product design.
- Provide relevant contextual data about active support issues so users better understand a service status.

Scalability: A key development goal for The Dashboard 2.0 is to design a system elastic in nature that can grow with demand and scale horizontally during times of high use and back down to a minimal footprint when demand is low. We accomplish scalability by implementing

standardized alerting, logging standardization, component decoupling and self-service. The introduction of a self-service model in conjunction with Role-Based Access Control (RBAC) will enable service owners to act on their respective services and create individualized alerts.

Reliability: We separate core components into a microservices architecture enabling the testing of individual pieces of The Dashboard much faster than is possible today. The introduction of a SDLC framework achieves software engineering maturity as we introduce new capabilities and features.

We have progressively introduced **Software Quality** and left-shifted across the SDLC, infusing quality assurance (QA) best practices into our SDLC framework.

We use the following approaches to provide high reliability:

- Development methodologies: Test Driven Development (TDD), Behavioral-Driven Development (BDD) and implementing failure recovery automation.
- The Defect Management Process will include defect workflows, bug and issue triaging, software QA metrics and reporting, automated testing framework and processes, regression testing and end-to-end testing.

Extensibility: We enhance The Dashboard's ability to extend into new areas to continue to build on the "single pane of glass" theme and onboarding new services. Options AGCY can choose from in this approach include an AWS cost dashboard, security dashboard and SSL certificate monitoring.

System Stability and QA: Our software QA process will introduce necessary tools to continue to improve system stability. Software best practices such as code reviews, pair programming and version control schemes provide mature code development.

Key benefits include:

- Platform extensibility and upgrades to the tech stack are made much easier and more flexible as components are decoupled.
- Features will be enabled or disabled using feature flags, scoped via canary deployments, to groups of users at a time.
- The N-tier architecture effort will build in added maturity, growth and long-term sustainability by following the AWS well architected framework.

3.3 The Dashboard: Core Product Infrastructure Capabilities

As The Dashboard exists within the AWS platform, we will continue to leverage cloud native technologies to power our system (**Table 1**).

Table 1: Infrastructure Resources

RESOURCE	USAGE
AWS EKS	Container orchestration tool that allows us to launch containers that run our backend and front-end technologies.
AWS Fargate	On-demand compute orchestration is used in conjunction with AWS EKS to scale containers up and down.
AWS ELB	Load balancing across high-availability front and back-end compute hosts.
AWS Kinesis	The messaging bus is used as an event hub; it passes information such as machine learning data to consumers, which manipulate and load to where necessary.
AWS Lambda	A stateless microservice, which we use to complete tasks such as alerting/notifications and data cleaning.
AWS DynamoDB	A nonrelational database we use to function as a metadata store, service definitions, metric definitions and service summaries.
AWS Sagemaker	Machine learning training and inference environment.
AWS Parameter Store	A distributed parameter store holds encrypted secrets such as API keys, login tokens and necessary parameters retrieved during execution run time activities.
Amazon Timestream	Serverless time series database used to store raw metrics.
Aws s3	Blob data store used to house unstructured and semi-structured data.
AWS Lake formation	Fully managed service used to build, manage and secure data lakes.
AWS Athena	A data warehouse interface technology that enables interfacing with a data lake and/or data warehouse.

Maintenance Efficiencies: To reduce maintenance activities, The Dashboard utilizes serverless compute technologies such as AWS Lambda and AWS DynamoDB. We use AWS Lambda to complete alerting tasks by reading from AWS Kinesis, a messaging event bus.

Infrastructure Monitoring: We accomplish infrastructure monitoring of our own service with AWS CloudWatch and send application-level logs to Splunk. We plan to manage The Dashboard user access control via HARP, authenticating users via an industry standard SAML 2.0 protocol.

4. THE DASHBOARD: ARTIFICIAL INTELLIGENCE (AI) AND MACHINE LEARNING (ML)

4.1 Introduction

The Dashboard and its AI/ML capabilities are built based upon the fundamentals of AIOPS, the natural evolution of application of Machine Learning to IT Operations, with the intent to improve IT operations holistically. This approach provides persona-driven data analytics via both historical IT operational data, streaming analytics at point of ingestion, and application telemetry.

> According to Gartner, **"There is no future of IT operations that does not include AIOPS. This is due to the rapid growth in data volumes and pace of change, exemplified by rate of application delivery and event-driven business models that cannot wait on humans to derive insights."**

Rapid growth in volume of data and complexity of infrastructure systems necessitates a forward-thinking, scalable AIOPS-based approach to enterprise system monitoring systems. Digital transformation is underway in many organizations that add complexity to existing IT systems, needing to leverage AI and ML capabilities beyond just software development, to provide a layer of 'human like' intelligence in the entire IT lifecycle across disparate systems not easily performed by humans. Ventech Solutions builds an AIOPS approach with a clear vision into the future of the possibilities of the application of AI and ML. Ventech Solutions' vision for The Dashboard is to provide that layer of visualization into an AI-enabled QNET environment.

Benefits to AGCY: The eventual transition into a proactive issue response culture from a reactive one helps stay ahead of the demands placed by digital transformation on IT systems and infrastructure today. This proactive culture will reduce operational downtime, allow for better cost efficiency and improve overall service to ADOs and end consumers. The Dashboard with AI/ML capabilities will also provide the benefit of analytical insights leading to improved infrastructure efficiency and lower costs for services consumed.

4.2 The AI/ML Learning Curve

More than 70% of AI/ML projects are canceled due to what businesses perceive as lack of Return On Investment (ROI) and the slow time to realize value, according to leading industry sources such as Gartner, Forbes, Fortune, and others.

> ***With the current AI/ML team, Ventech Solutions** has overcome that learning curve, as demonstrated by 16 custom ML models developed within the last six months.*

4.2.1 Challenges Overcome

Environment and team: AI/ML capabilities are not easily accomplished, especially in standard systems development, due to several challenges, including but not limited to finding high-quality talent with data scientists and ML engineers, the ramp-up time it takes to understand the challenges within the data sciences and the specific AI/ML environment for the capability. “Time to realize value” increases the time before any significant gains are to be realized by the new team, and these transitions are major contributors to AI/ML project cancellations.

Technical: Deriving analytical insights is challenging, given the data quality issues. Considerations contributing to these challenges are complexity, data integrity and quantity, data seasonality and stability, the complexity of algorithms and performance, and machine learning models’ interpretations.

4.3 AI/ML Business Capabilities

Our near-term enablement of data analytics within The Dashboard provides machine learning-based predictive analytical and anomaly detection capabilities to eight of the 66 services (to date) currently onboarded to The Dashboard.

We provide the following capabilities, with the associated benefits:

- **Predictive analytical capabilities**:

  - We provide a window into the future of health of key service KPIs from two to up to 30 minutes (depending on data quality for machine learning).
  - Service owners are provided an early warning based on the KPIs they deem are important, enabling service owners to take proactive mitigation actions before waiting until their services were down.

- **Anomaly detection capabilities**:
  - Our advanced ML capabilities detect outliers in the data and abnormal patterns in historical data that point to anomalous conditions impacting service uptime.
  - Service owners are alerted, allowing them to perform root cause analysis for mitigation of both current and potential future problems.

- **Success stories:** Recent success stories with this capability have been proven by the Ventech Solutions team when anomalies in the Barracuda service were detected by our machine learning models preventing system degradation. The in-house monitoring products the service team used did not detect those anomalies.

- **Correlation identification:** Part of the AI/ML capabilities planned includes the identification of correlation between services by discovering patterns in the data that might provide these insights. This identification will provide capabilities such as the potential impact of one service going down to the other services.

- **Model interpretation/**explain**-ability:**
  - One of the challenges when working with machine learning models is interpretability. Current ML models are a black box, and the ability to understand the model prediction outputs—is a key benefit. This model also builds user's confidence with machine learning.
  - The Ventech team is researching the use of cutting-edge algorithms to explain models' prediction to help users understand the model outputs, including benefits such as what KPI might fail with a given IT service or where the anomaly is.

- **Humans in the loop:** As part of the AI maturity, we involve the human element within the AI/ML lifecycle for continuous improvement and effective and meaningful AI implementations.

- **Data capabilities:** Currently, Ventech Solutions' capabilities in this space are with multivariate time series datasets. Our AI/ML team has the capability to grow forward, working with a variety of datasets, including big data as well as structured, semi and unstructured datasets.
- **Anomaly detection** and prediction capabilities are enabled where we have quality data for ML use. ML model improvements and continuously enhancing model accuracy are part of our objectives.

4.4 AI/ML Technical Capabilities

4.4.1 Machine Learning Architecture

We currently use ML model architecture called LSTM (Long Short-Term Memory) neural network; it is a form of a recurrent neural network used in deep learning methods.

4.4.2 Model Performance Metrics

We have achieved the following model accuracy and precision, with model performance metrics around < 8% Mean Average Prediction Error (MAPE) for 5-minute models, and < 15% MAPE for 30-minute prediction models. For anomaly detection models we have model performance metrics of F1 score of > 0.85.

4.4.3 Data Engineering Considerations

Data at the right quality level, suitable for machine learning consumption, is key to success. There are considerations we have factored in, including, but not limited to: data integrity and quantity, seasonality, stability, algorithm complexity, performance and ML model interpretation.

4.4.4 Machine Learning Models In Use Today

The Ventech Solutions team has evaluated the following ML models in this environment:

- **Prediction:** A wide range of machine learning algorithms are in use today:
    - **Regression Models:** We evaluated Autoregressive Integrated Moving Average (ARIMA) models to predict univariate KPIs for five minutes of prediction; Vector Autoregressive Models (VAR) are used to predict multivariate KPI values.

  - **Deep Learning Models:** LSTM networks predict up to 30 minutes into the future. This model is the final architecture selected due to superior performance over ARIMA and VAR.
- **Anomaly Detection:** Unsupervised and self-supervised machine learning methods; we evaluated both Extended Isolation Forest and Deep Autoencoder models, with a final selection of the LSTM Autoencoder architecture adopted for use.
- **Ongoing ML Research:** We continue researching additional, more cutting-edge ML models to continue to improve and optimize the value from The Dashboard.

4.4.5 Machine Learning Model Maturity and Improvement

As a part of continuous maturity and improvement, our models are designed to handle data drift. We continuously monitor anomaly detection model reconstruction errors and trigger models retrain when needed. We estimate data baselines and adjust model thresholds accordingly.

4.4.6 Machine Learning Operations (MlOPS) - Platform Capabilities

- **ML Development:** The core activity during this ML development phase is primarily experimentation. Our ML engineers perform the following activities: Prototype model architectures, data preprocessing, feature engineering, model training, model evaluation, hyperparameter tuning and optimization. The primary output of this process is a formalized training procedure that includes the following: Data preprocessing outputs, model architecture and model training settings.
- **Continuous Training:** The continuous training pipeline is executed repeatedly based on retraining triggers and produces a model as output. The model is retrained as new data becomes available or if model performance decay is detected. Should the ML pipeline produce a successful model candidate, that candidate is then tracked by the model management process as a registered model. The models are retrained after deployment to production when model decay occurs or is detected.
- **Platform Optimization:** The platform uses GPU instances with an optimized infrastructure to reduce training time from hours to minutes. This approach has shown us it boosts team productivity up to 10 times. ML managers annotate and review the

registered model and approve it for release. The model is then deployed to a production environment. We built a custom DevSecOps pipeline for incremental delivery.

- **Serverless Inference:** Serverless Inference is integrated with our MLOPS pipelines to streamline ML workflows. We incorporate a serverless endpoint to host the model that has been registered with the Model Registry. Serverless endpoints automatically launch compute resources and scale them up or down depending on traffic, eliminating the need to choose instance types or manage scaling policies. With a pay-per-use model, the Serverless Inference is a cost-effective option that helps us to reduce cloud costs from thousands to hundreds per month.
- **Continuous Monitoring:** The continuous monitoring process monitors the model for effectiveness. The objective of effective performance monitoring is detecting model decay and data and concept drift. The model deployment process monitors for efficiency metrics: Latency, throughput, hardware resource utilization and execution errors.

4.4.7 Future Work In AI and Other Advanced Algorithms

The Ventech Solutions team is exploring advanced AI techniques to better predict and interpret KPIs and identify anomalies:

- Natural Language Processing (NLP) technologies:
  - Several algorithms are planned for future AI/ML research and applications: Latent Dirichlet Allocation (LDA), Hugging Face Transformers libraries, SHapley Additive exPlanations (SHAP), Bidirectional Encoder Representations from Transformers (BERT), RoBERTa and other pre-trained Sentence Transformer models (SBERT). These can be used to analyze text conversations and application log messages, instead of numeric metrics, as textual messages usually provide contextual information into the systems status.
- Deep Reinforcement Learning (DRL) is another technique on the team roadmap to tackle dynamic data and service changes. ML models will make optimal decisions (anomaly or normal) based on continuous learning, thus minimize the risk of model drift.

- **Table 2** highlights ongoing Proof of Concept (POC) initiatives that are part of the Dashboard Program.

Table 2: AI Capabilities in the Dashboard

AI Model	Model Architecture
Uptime prediction	Time series LSTM
Anomaly detection	Time series LSTM Autoencoder

Advanced AI	Algorithm and Model Architecture
Time series prediction	ARIMA, VAR
Time series prediction	Temporal Convolutional Network (TCN)
Black Box Models explain-ability	SHAP
Text sentiment analysis	CBOW, BERT
Topic modelling	LDA, BERT
Dependency analysis	Granger Causality
Predictive anomaly detection	Darts forecasting anomaly

4.5 The Dashboard AI/ML Maturity (Long-term Vision)

Ventech Solutions' vision is to enable AI within The Dashboard to leverage gains from AI/ML capabilities. We are planning the following activities:

- NLP capabilities for IT service status use and other business outcomes related to QNET environment health that can benefit from consumption of analytical insights.
- Analyze Chatbots utilizing sentiment analysis.
- Benefit: This capability has the potential to lead to discovery of real-time issues other ML approaches are unable to detect.

- Other analytical capabilities such as:
  - The ability to gain analytical insights on text-based data in the environment.
  - Slack and Confluence data mining with techniques such as 'Bag Of Words' / 'Continuous Bag Of Words' models (BOW/CBOW).
  - Topic identification for analytical use cases.
  - Automated response with chatbot and voice activation.
  - Sentiment analysis on Slack message channels supporting system health status and confluence data mining for analytical insights.
- Determine causation and correlation between services:
- Discover correlation between services that provide potential impact of one service going down to correlated services being impacted and degraded.
- Web scraping for data analytics
- Develop web scraping of data (both the environment and the WWW) for data that can be used for sentiment analysis, with the benefit of improving machine learning outcomes and deriving additional analytical insights.
- AI and RPA integration:
  - Leverage predictive capabilities to spawn any automated (robotic) process to go execute 'something.' Example: Automated server upgrades and patches or enhancing infrastructure or system capacity.
  - Evaluate one pilot use case with an ADO.
- Evaluate/Research Intelligent alert escalation and routing:
  - Use of AI capabilities to alert users of potential threats to environment.
  - Evaluate one pilot use case.
- Expand Historical analytics development:
  - Evaluate one pilot use case.
- Evaluate AI assisted triaging: Evaluate ways to use AI with human-driven triaging, such as incident remediation.
- Research ways to use AI/ML to build resiliency in the environment. Pilot with a partner for one use case prior to expanding.
- Research ways to build application predictive capabilities in the environment and detect seasonal application trends. Pilot with a partner for one use case prior to expanding.


Ventech
Solutions