Building incident avoidance using digital cockpit– design principles

wipro

CIOs today are challenged with meeting business expectations to perform on par with consumer-facing technology leaders like Facebook and Netflix—with always available and zero latency performance. Of course, these expectations are not aligned with the realities of managing complex hybrid enterprise systems that were not necessarily designed to deliver this type of performance.

The **new generation hybrid enterprise** has to deliver the functionality of, and operate as, a **real time infrastructure.** It is also built on microservices and cloud native platforms and is inherently designed to fail. This is what the IT Operations is tasked with supporting.

Digital enterprises that use such hybrid environments and must respond in real time need to achieve incident avoidance, rather than incident management.

But most operation command centers today struggle to keep up with the required pace. Very few operations have been built from the ground up, which is why new age platforms like Facebook and Netflix are in a different league altogether. Today Severity 1 incident resolution times still run for hours with a large group of SMEs on bridge calls, with no root cause to be found and each merely looking to avoid blame. **Lack of a proper CMDB, insufficient logging, or lack of application discovery and dependency mapping are the culprits.** Every IT professional recognizes these failings, yet projects to correct them never seem to pass the capital appropriation business case hurdles.

A new approach is needed. We propose a set of design principles for an operationally simple, scalable, and focused solution that brings IT and business process owners to a common understanding.

Design principle #1: Simplicity is key—focus on progress, not perfection.

To ensure simplicity, start with a clear understanding of the services, transactions and functions. The conventional wisdom stipulates that we start with a robust and expensive infrastructure and application monitoring that will gather significant amounts of system and application data. However, when we cast the net wide, we fail to consider that this data is in itself irrelevant until and unless there is a disruption in a business service.

Similarly, striving for progress and not perfection stipulates that we accelerate our rate of learning rather than seek to perfectly maintain every system or subsystem.

The first design principle is to **capture the business service heartbeat.** In a clinical sense, we seek to uncover the smallest atomic transaction/service component. Building our monitoring on this foundation allows us to circumvent the complex and often expensive CMDB and ADDM projects. Depending on the industry, this heartbeat might be customer orders, service requests or vendor purchase orders. The other business processes are subordinated to this heartbeat.

Design principle #2: Shift from monitoring events to raw wire data (from thresholds to live streamed data).

As operations experts, we spend much of our careers delivering complex systems built on the installation and configuration of multiple tools and application instrumentation (logging and defining policies about data retention for each tool). We then struggle to perform the real magic of root cause identification and analysis to configure alerts based on defined thresholds.

Fundamentally, what we want to measure is dynamic, but how we measure it remains static. This needs to change.

Our second design principle defines the way to change this and looks at streaming analytics-based approaches as shown in the figure below.

We should take a heartbeat of critical business processes, and provide real time analytics that will change the way our knowledge workers approach problem solving. **Our static, threshold-based models have created a generation of IT professionals who wait for the incident rather than anticipate problems. This premise needs to change.**

Digital enterprises that use such hybrid environments and must respond in real time need to achieve incident avoidance, rather than incident management.

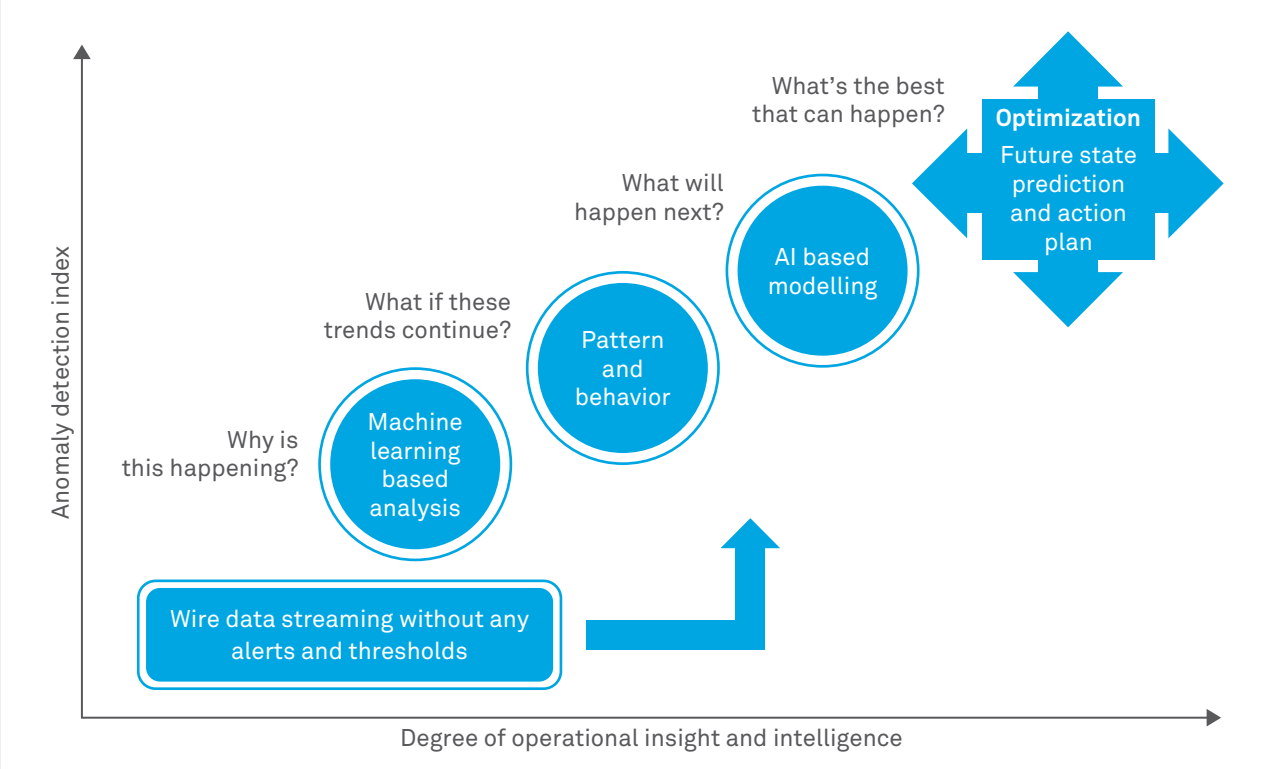

Fig 1: Using streaming analytics to build real time models and patterns

Where evaluation of process health is truly continuous, the assessment and reaction must be near-instantaneous. This health needs to be visually depicted in a single pane with the highest degree of insight and intelligence, and in context, thus helping avoid an incident.

Design principle #3: Detect anomalies in the business process based on KPIs learnt, without waiting for underlying infrastructure failures.

To rethink system management, we must redefine how we uncover defects. This approach traditionally starts with creating a project for service maps. Instead, we capture the **most atomic transaction and then start learning its nature.** It is this nature that embodies the collection of underlying systems as well as user behavior and even seasonality. From this nature we build behavior patterns and then models, using which we can predict and prevent.

Imagine if we used the flow of service requests from a collection of global customers as our heartbeat. As we watch this stream of data, we get a sense of when specific regions start their business day, how our IoT collection batches operate, and even when the fulfilment teams

change shifts. Everything we need to operate our business is embedded in this stream of data. The question is how to uncover it.

The next design principle explains how we can achieve this using data discovery tools and a machine learning platform.

There will always be a defined set of sources logging business processes. By studying some past historical outages and tracking when the "symptoms" started, one can determine where to start and what key words to search for, to develop **simple, expression-based data discovery.**

Once we have started, we can then continually iterate using machine learning. Since the system will be in use, it will have a more relevant and practical feedback mechanism. The flow chart below suggests a simple method to begin, establishing a platform from which to learn and progress.

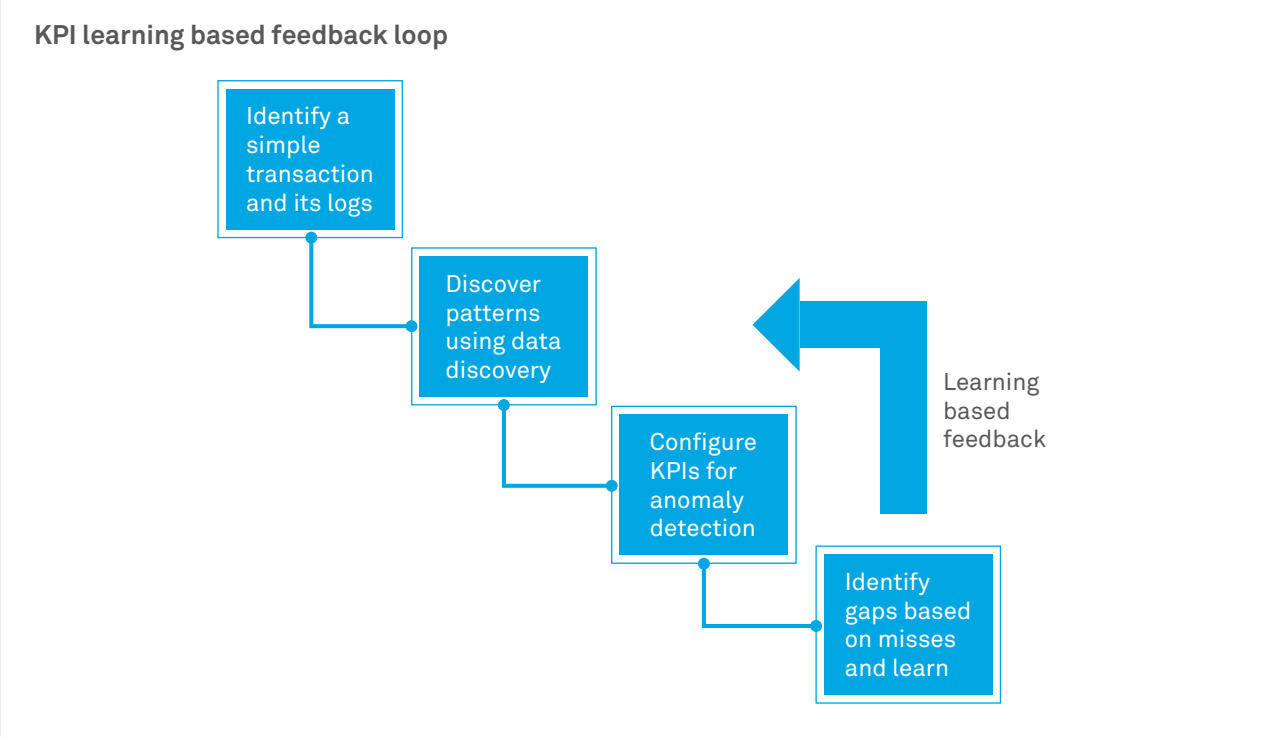

Fig 2: Using machine learning for identifying KPIs and enhance through feedback loop

Thanks to machine learning platforms combined with Big Data analytics, finding KPIs and a baseline is a much easier task today. A flexible machine learning platform allows us to start experimenting with the existing historical data available, before investing in configuring the real-time dashboards. In a small system, the number of components involved are also small and thus the complexity is low. Once proven as a base case for a small system, one can proceed with setting up a platform using that as a template.

Design principle #4: Start with the digital cockpit visualization to guide the enterprise.

Every successful endeavor started with some kind of visualization that sounded impossible to implement. Without taking any one particular existing model, use the simple transaction chosen for the atomic heartbeat, and with the KPIs acquired in consideration, "draw" what you would like to have in a **real time dashboard.** The key attributes of that dashboard must show transaction/service health in real time. Anomaly needs will be detected at the business process level.

This design principle describes how to build such a visualization by following a bottom up strategy, as shown in the figure below:

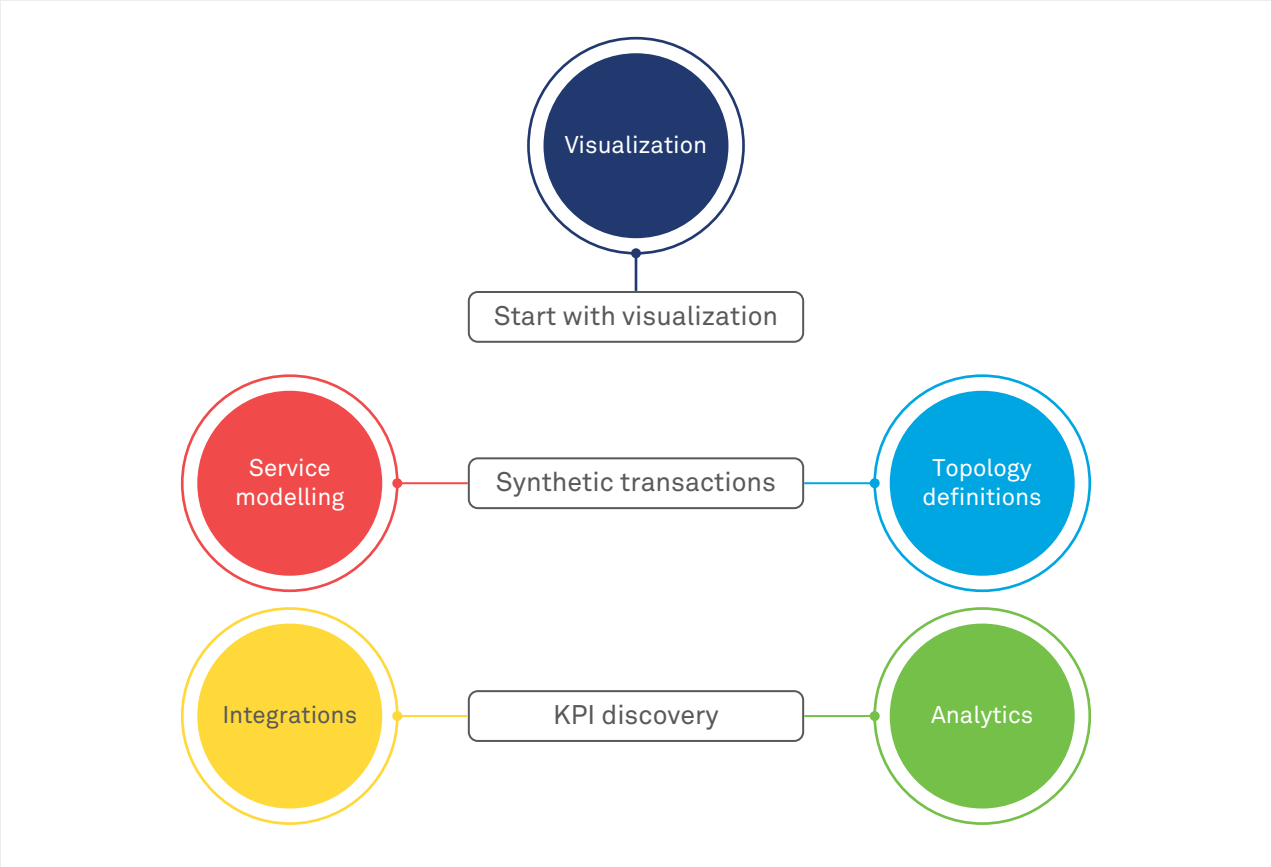

Fig 3: Start with visualization first

Machine learning, Big Data analytics platforms and several enterprise tools and platforms that already exist in the enterprise can help create this visualization.

For most enterprises there is no need for a new tool or platform to build a visualization. An existing tool or platform can be used.

Incident avoidance, not incident management

Integrated IT operations is not a goal; it is the means by which digital enterprises can be managed in real-time. The path to real-time management is through service capacity guarantees, performance and availability management, using anomaly detection in the context of the business process to identify issues well before an incident occurs. In practice, start small and build an anomaly detection system for the heartbeat of the process that enables you to see both the forest and the trees.

About the author

Murthy Malapaka
Head of Transformation Services -
North America, Wipro Limited.

Murthy has 28 years of experience as a technology innovator and change agent. Over these years, he has assumed various technology leadership roles across application and infrastructure architecture domains, specializing in availability and reliability.

He has been providing consulting services to CIOs and CTOs in their journey from client server to on demand infrastructure services.

Murthy can be reached at
murthy.malapaka@wipro.com

Wipro Limited
Doddakannelli, Sarjapur Road,
Bangalore-560 035,
India

Tel: +91 (80) 2844 0011
Fax: +91 (80) 2844 0256
wipro.com

Wipro Limited (NYSE: WIT, BSE: 507685, NSE: WIPRO) is a leading global information technology, consulting and business process services company. We harness the power of cognitive computing, hyper-automation, robotics, cloud, analytics and emerging technologies to help our clients adapt to the digital world and make them successful. A company recognized globally for its comprehensive portfolio of services, strong commitment to sustainability and good corporate citizenship, we have over 160,000 dedicated employees serving clients across six continents. Together, we discover ideas and connect the dots to build a better and a bold new future.

For more information, please write to us at **info@wipro.com**

IND/B&T/NOV2018-APR2019