

WHITE PAPER

Infrastructure AI Operations at Hitachi Vantara

AI and Machine Learning Initiatives
for Data Center Modernization

By Hitachi Vantara
April 2020

Contents

- Executive Summary** 3
- Data Center Operations Demand AI** 4
- Hitachi AI Offerings for the Data Center: Today** 4
- Ops Center Automator and Machine Learning** 4
- Ops Center Analyzer and Machine Learning** 5
- Hitachi AI Offerings for the Data Center: Tomorrow** 6

Executive Summary

As the diversity of systems and software used to drive business operations has increased, so has the amount of time data center teams spend managing those systems. In many cases, data center teams must make hard choices around how much time they spend managing production systems versus planning and implementing new initiatives.

To help free up staff to research, architect and deploy new solutions, data center leaders have traditionally leveraged monitoring software that can continuously check system faults and outages so that operations do not halt. This approach helps ensure service level agreements (SLAs) for uptime and data protection are met. However, it does not provide the real-time monitoring needed to ensure systems and software are optimized individually and as a holistic unit, including application, hypervisor, server, network and storage. All need to be adjusted for peak performance and uptime as well as maximum return on investment (ROI).

To get the most out of data center investments, make sure budgets are being effectively leveraged, and minimize the risk of outages, organizations are turning to increasingly powerful artificial intelligence (AI) offerings. With AI, leaders see the opportunity to focus more time on strategic initiatives while the data center monitors itself in real time.

In this paper, we explore how, where and why Hitachi Vantara is integrating AI and machine learning (ML) into data center management software. Our discussion includes a look at both current capabilities and areas of development where Hitachi sees opportunities to enable deeper insights into how data center operations, both individual and global, can be optimized.

Data Center Operations Demand AI

Managing the wide mix of systems and software that power digital business is no easy task. To get the most out of data center investments and minimize the risk of outages, systems and software must constantly be monitored: and not just at a device level. They must also be monitored in groups based on how they work together (for example, an application working with server, network and storage).

Monitoring and analyzing operations to maximize performance, forecast ongoing resource needs, and identify, then fix faults can easily consume the bulk of a data center team's time. And while this is important for day-to-day operations, it also pulls teams away from planning future initiatives.

To better balance where time is spent, infrastructure and operations leaders are turning to artificial intelligence offerings. With AI, leaders see the opportunity to focus more time on strategic initiatives while the data center monitors itself. In effect, this allows an organization to begin implementing an autonomous data center that intelligently predicts and prescribes changes that increase uptime and operational efficiency.

Hitachi AI Offerings for the Data Center: Today

Designing AI and ML capabilities is a major focus for Hitachi engineering teams. Across every area of the company, development teams are researching how to develop AI to make smarter, broader insights. In the data center, Hitachi's AI and ML efforts are currently focused on Hitachi Ops Center Automator and Hitachi Ops Center Advisor. These offerings are designed to predict, prescribe or execute actions based on two core approaches to machine learning:

- **Decision-tree learning.** AI examines the status of a device (for example, the storage array) and uses a series of logic operations to predict what will happen next and prescribe changes to configurations.
- **Association rule learning.** AI examines a set of devices (for example, server, virtual machine, storage and network) and correlates information from the devices to identify and prescribe changes to the resources so a desired outcome can be reached (for example, best deployment model, root cause analysis).

Combined, they help support organizations' journey to an autonomous data center, where day-to-day operations are managed by intelligent systems and software. How each learning method is executed depends on the individual product.

Ops Center Automator and Machine Learning

Ops Center Automator simplifies the process of deploying and managing data center infrastructure by automating common configuration and administrative tasks. Unlike other orchestration offerings, Automator was built around the idea that resource management cannot be based on "static" assignments (for example, if you need capacity, allocate it from array x). Instead, decisions must be made based on the "best" option at any given time (for example, if array x is overloaded, choose a different array to use).

By considering the current state of resources, Automator is able to improve resource utilization and optimize performance. To do this, Automator includes two key functions: smart storage provisioning and fabric-aware provisioning.

- **Smart storage provisioning:** With this function, the embedded AI analyzes storage telemetry data and prescribes the best available resource to use to meet a desired set of goals. To accomplish this, Ops Center Automator collects and analyzes the available capacity and busy rate across a set of storage resources. Automator then maps the storage array telemetry detail to the administrator's desired service level objectives and goes on to determine the most appropriate storage array and media type to use. This model is programmable so that the decision map and the outcomes can be adjusted.

– SUPPORTED DEVICES: Hitachi storage arrays.

- **Fabric-aware provisioning:** This function works in a similar fashion to determine the most appropriate data path and switch-zoning configuration to use when deploying a new application. Here, the embedded AI takes input on network zoning requirements, such as number of paths and multipathing requirements. Then, it analyzes the network topology between server and storage end points. Based on requirements, Ops Center Automator prescribes the best zoning configuration to implement.

- SUPPORTED DEVICES: Brocade and Cisco SAN switches.

Decision-making within Automator is increasingly linked to other offerings that have broader analytic capabilities or external management practices. This includes Ops Center Analyzer and third-party IT service management (ITSM) tools, such as ServiceNow.

Ops Center Analyzer and Machine Learning

Hitachi Ops Center Analyzer analyzes telemetry data from resources across the data path – hypervisor, server operating system, network and storage – to isolate, predict and prescribe changes needed to keep the data center running smoothly. Analyzer accomplishes this by leveraging a mix of decision-tree and associated rule-learning models, depending on the desired outcome.

- **Capacity (Budget) Forecasting:** With this function, Ops Center Analyzer analyzes capacity utilization trends for a storage array over the course of a defined time period. Analyzer learns the pattern of how the resource is being used and designs a forecast for when additional capacity will be needed to keep operations running. Analyzer allows the level of standard deviation to be adjusted so administrators can plot “what if” scenarios based on changing workload capacity requirements.

- SUPPORTED DEVICES: Hitachi storage arrays today.

- **Performance Forecasting:** This more advanced forecasting function analyzes performance telemetry to determine when a resource will hit peak performance. Here, Ops Center Analyzer uses data acquired from the critical resource, including periodic patterns (for example, weekly spikes for batch jobs), and learns how the resource is behaving. Analyzer then predicts how the resource will behave in the future and when it will reach a maximum performance threshold. Analyzer allows the level of standard deviation to be adjusted so administrators can plot “what if” scenarios based on increasing or decreasing performance needs.

Analyzer performance forecasting is different from other offerings because it also examines dependencies: resources that influence the performance profile (for example, cache, ports, storage and so forth). By considering other resources in the critical path (for example, whether a network bottleneck will cause the VM to hit peak performance) staff can more accurately determine what needs to be upgraded. This approach prevents budgetary waste on devices that may not be at peak usage.

- SUPPORTED DEVICES: Previously listed storage arrays, Brocade and Cisco SAN switches, Linux and Microsoft Windows operating systems, Microsoft Hyper-V and VMware hypervisors.

- **Performance Threshold Management:** This function enables Ops Center Analyzer to predict and adjust the performance settings of various resources in order to protect customer experiences. Analyzer begins by monitoring response times and/or IOPS for a resource. Telemetry data is compiled and outliers are removed as part of the learning process to understand the normal operating behavior of the resource. A normal operating band is calculated and updated over time to find anomalies in performance.

- SUPPORTED DEVICES: Previously listed storage arrays, Brocade and Cisco SAN switches, Linux and Windows operating systems, Hyper-V and VMware hypervisors.

- **Anomaly Detection and Root Cause Analysis:** With this function, Ops Center Analyzer monitors the operating ranges of resources in a data path for anomalous behaviors (for example, performance spikes). Based on defined SLAs and behavior over time, Analyzer decides if an error condition has been reached.

Once Analyzer identifies an anomaly, it analyzes the device with the issue as well as dependent resources. Associated rule learning takes performance telemetry from across resources as well as detail on recent configuration changes to prescribe resolution. This could include adjusting configuration settings or modifying quality of service (QoS) settings to eliminate noisy neighbor impacts.

If Analyzer prescribes a change to QoS settings on storage, staff can automate updates to the storage array configuration directly from the Analyzer interface.

- **SUPPORTED DEVICES:** Previously listed storage arrays, Brocade and Cisco SAN switches, Linux and Windows operating systems, Hyper-V and VMware hypervisors.

In each of the scenarios outlined, Analyzer is able to make decisions locally without engaging a remote intelligence. This capability is extremely important for security-conscious and government facilities that do not permit outbound traffic or communication of infrastructure operational details.

Hitachi AI Offerings for the Data Center: Tomorrow

As Hitachi moves forward to help accelerate the move to autonomous data center operations, we are making investments in three areas:

1. Expanded infrastructure device support.
2. More advanced learning models.
3. Collaborative AI.

While Hitachi already supports a superior number of Hitachi Vantara and third-party infrastructure devices, we recognize that this listing does not provide complete coverage of all the systems and software that power digital business. For this reason, we are actively developing support for additional vendor systems as well as key applications. Our goal is to enable deeper, more accurate insights and to eliminate analytics “silos” that exist in other vendors’ products that only consider their own systems.

In parallel to this effort to expand system and software support, Hitachi teams are working on more advanced machine learning models that will enable us to predict and prescribe a broad range of events. Development efforts are already underway on building unsupervised learning models and natural language processing, to both recognize different dialects and identify key words within logs and data streams. Such models are intended to improve the speed at which our AI matures and the benefit it provides to organizations, worldwide.

At Hitachi, we recognize that the speed of innovation is determined by the ability to collect and leverage data. The same is true when automating data center operations. By connecting different information sources, even different AI, we see the opportunity to provide exponentially better predictions and prescriptions: We see where we can make data center operations run more efficiently and with less downtime.

To enable this, all of our data management offerings are designed with the ability to share information across API interfaces and work together. Already, work is underway to link Ops Center Analyzer and Ops Center Automator with various offerings, including:

- **Hitachi Remote Ops monitoring system:** The Hitachi Vantara support system. AI in Remote Ops leverages a 10PB+ data lake built from details across 13,000 customers and 64,000 systems. Roughly 40 trillion data points per day are monitored, to identify and prescribe system level changes so organizations avoid prolonged outages.
- **Hitachi Ops Center Protector:** This software monitors and manages data protection for a wide range of systems. Integration with Ops Center Automator enables automated setup of data protection policies on systems and informs Ops Center Protector how to monitor and “fix” policies that are not operating as expected.

- **Pentaho Solutions:** The Pentaho analytic framework allows predictive analysis across data centers and across the broader range of information sources that govern data center operations. This includes cameras, power and cooling, badge access, weather information, social and news feeds (to determine human impact on operations).
- **Third-Party Sources:** Via API, Hitachi is looking to integrate third-party offerings. This approach could include analytic sources, such as application monitoring software, for deeper insights and more precise decision-making. Or, it could employ external orchestration frameworks to leverage our broad range of insights, to improve their own decision-making and better visualize corporate operations.
- **Autonomous Systems:** Every Hitachi platform is defined, at its core, to collect information about its operating state and make predictions based on that telemetry data. That capability will be expanded over time so that our platforms can be self-healing and work with other systems, such as network devices. It will enable them to make tactical decisions on how to optimize operations (for example, work around faults or bottlenecks) and meet organizational goals set by an overarching framework.

Going in depth on each of these items is outside the scope of this paper. Instead, we've presented Hitachi Vantara's AI focus. This focus spans many dimensions that are based on collaboration and an open framework that allows organizations to use AI and ML to free staff to focus on more strategic endeavors while the data center manages itself.