A HURWITZ WHITE PAPER

Innovation: Add Predictability to an Unpredictable World

Improve Visibility and Control of Your Telecom Network

Judith Hurwitz
President and CEO

HURWITZ & ASSOCIATES
Insight to Action

Sponsored by Hitachi Data Systems

Introduction

It is all too easy to use the word innovation to explain how companies can differentiate themselves from each other in a world where commoditization is becoming commonplace. However, it is clear that companies that thrive are the ones that are able to provide a level of assurance and predictability for customers. It is clear that to be successful, organizations need a foundation of data to be able to manage, monitor, anticipate and respond to customer requirements.

Each type of business has requirements to harness the myriad of data sources in a proactive manner to create sustainable value. In this paper, we will examine the challenges and requirements of the telecommunications industry and how the Hitachi Data Systems approach to big data management is supporting the changing requirements in the industry. In this paper, we focus on the complex market for network operations centers in telecommunications.

The Challenges of Telecommunications Providers

Telecommunications providers are challenged with the requirement to offer competitive and reliable connectivity services for individual subscribers, enterprise customers and even other service providers. Increasingly, the expectations of constituents are changing as new real-time services are introduced. Whether the demand comes from the consumer accessing a video in real time or from a backhaul service provider that offers services to other providers, consistency and predictability are required. In a highly competitive market, lack of a proven service level agreement (SLA) can result in loss of customers and the inability to expand to new initiatives. In addition, service providers are at risk of incurring penalties if they do not meet service level objectives (SLOs) and SLAs.

In a highly competitive market, lack of a proven Service Level Agreement (SLA) can result in loss of customers and the inability to expand to new initiatives.

The Complexities of Controlling Network Traffic

One of the biggest obstacles facing telecommunication companies is having the right level of visibility and control to be able to anticipate and prevent outages or network slowdowns that can impact customer satisfaction. Even more significant is that these telecommunications companies can suffer a loss of data through packet loss.

Typically, a network operations engineer will use a performance monitoring application to manage traffic. It is quite common that network traffic and performance data is collected and reported in 5-15 minute intervals from individual network elements, and sometimes up to one-minute intervals with network probes. In the past, this was an acceptable way to manage traffic. However, with real-time services becoming the norm, there is a need to refresh traffic data more frequently. For example, if a customer has bursts of traffic, it could happen within 30 millisecond intervals. While traffic averages may appear to be smooth and compliant, there may be serious flaws that will go unnoticed.

HURWITZ & ASSOCIATES
Insight to Action

The enterprise customer will notice unacceptable latency and data loss in their applications. However, when that customer contacts the support center, the network operations center is unable to detect the problem.

In complex environments, these average refresh rates are no longer acceptable. Most service providers require refresh times of below one second because of the guarantees they must provide to customers. For example, the customer watching a popular program on Netflix will not be happy with a 10- or 20-second delay in transmission. Commercial backhaul providers must be able to offer 4G, real-time SLAs to customers for smartphone connectivity. These providers must demonstrate how they will deliver traffic accordingly, and maintain end-to-end delay within the 10s of milliseconds. If not, they will face stiff penalties.

Having the ability to manage and analyze data in near real time is the heart of the way that telecommunications providers can prevent outages and meet customer requirements. The typical way a network operator determines the root cause of network latency is to analyze systems logs, which contain massive amounts of highly technical data that is extremely hard to correlate. Even with newer big data analytics tools, network operators need answers in real time. Being able to predict or determine quickly where the specific problem resides is critical in assisting the client, even when the problem resides within the corporate network.

Having the ability to manage and analyze data in near real time is the heart of the way that telecommunications providers can prevent outages and meet customer requirements.

The Role of Big Data Analytics

Collecting and storing data related to controlling network performance can be complicated. As we mentioned earlier, it is typical for service organizations to be able to capture snapshots of performance metrics so that an average can be calculated. It is often the case that the performance dashboard does not catch problems related to bursts where service delivery is uneven.

To be truly effective in an always-on world of communications, organizations must capture and analyze the streams of data in real time so that subsecond responses can be identified. This type of fast response will make it possible to correct problems before they impact performance during an unanticipated incident. However, to truly maintain the level of performance over the long term requires a combination of real-time data and historical data. In this way, it is possible to benchmark current performance against past performance. If a company is able to determine the best SLA that the customer expects, they can continually collect data and compare current data against expectations. Being able to do correlation analysis between real-time streaming data and historical data can enable organizations to accurately track performance.

Predictive and Prescriptive Analysis

While it is important to understand the immediate state of the network, it is even more important to ensure that the network can respond and change as user requirements and expectations change. This is especially important in a highly competitive market where emerging types of content and applications can

HURWITZ & ASSOCIATES
Insight to Action

suddenly change the performance of the network without warning. In addition, corporations are constantly changing physical regions, whether through mergers and acquisitions or through planned growth. As a company grows, it may be necessary to suddenly add several thousand users in a single month. Many companies are also faced with employees who work remotely across the globe.

Service providers that have the ability to predict impacts to the network based on real information as well as predictive analytics are in a better position to gain an advantage in contrast to providers who can only assume potential impacts. The burden is often on the client to provide assumptions, historical data and requirements for future performance. Many of these service providers are offering premium services to customers who require a higher level of network availability and performance to support their applications and customers. In fact, one of the main reasons that customers select premium services is to ensure that their online services for their customers and partners are reliable. Without a strong service and operational performance guarantee, companies risk losing customers or being able to support their own employees.

Certain techniques can help move the process forward. For example, it is possible to set threshold rules that define the acceptable service level. These rules can be nuanced to take into account thresholds that are defined for different types of customers and the services they have purchased. Thresholds can be monitored through probes and sensors deployed at the edge of the network as collection points.

But it is not enough to simply be able to monitor the thresholds because there are invariably outside influences on traffic patterns that cannot be anticipated. A service provider may be negatively impacted if there is a surge in traffic because a new television program debuts online. However, if that service provider has been monitoring social media and can correlate that information with network data, they could be better prepared to manage customer expectations. All of this requires a massive amount of data to make it possible to discern patterns and trends.

Another important ingredient in being able to both anticipate and make proactive action is the ability to use machine learning to allow the network services to morph and change based on the data that is being collected over time. This capability can allow an organization to measure new thresholds that have suddenly become relevant. An organization may estimate the right threshold for a network, but, as conditions change, this estimate may no longer be accurate. Machine learning, however, can determine changes and anomalies over time and dynamically change the threshold limits based on actual performance. This ability to be able to change thresholds dynamically is important in an era where new services are being offered and changed constantly. Using learning algorithms, the network behavior can change and the system can be orchestrated to be more prescriptive. Winners are those companies that can take action that meets changing customer expectations and requirements.

Using learning algorithms, the network behavior can change and the system can be orchestrated to be more prescriptive. Winners are those companies that can take action that meets changing customer expectations and requirements.

HURWITZ & ASSOCIATES
Insight to Action

Consider the Outcomes

All of this data and analysis is only useful in the context of predicting and changing outcomes. Many companies today are able to analyze data and understand what is happening and why a threshold is not effective, but they are often unable to determine what actions to take to impact outcomes. What is required is the marriage of data collection and analysis in real time combined with the ability to put those results in context with past performance. Successful companies are able to use this analytical approach to proactively partner with clients. Customers want to be able to have the assurance that their networks will scale and perform well as their requirements change.

Hitachi Data Systems (HDS) Big Data Framework

Hitachi has been creating solutions that address these types of big data problems for many years. To increase its viability in the market, the company is codifying its best practices into a solutions framework. This framework will be the foundation for a portfolio of end-to-end solutions in vertical markets, including telecom, healthcare, oil and gas, and others. The framework includes Hitachi intellectual property as well as APIs that link those services to 3rd-party offerings and open source platforms. Deployment is supported both on-premises and in the cloud. The framework incorporates services at a number of levels that enable both ingest of real-time streaming and historical data.

Hitachi has been creating solutions that address these types of big data problems for many years. To increase its viability in the market, the company is codifying its best practices into a solutions framework.

Figure 1. Hitachi Data Systems Big Data Framework

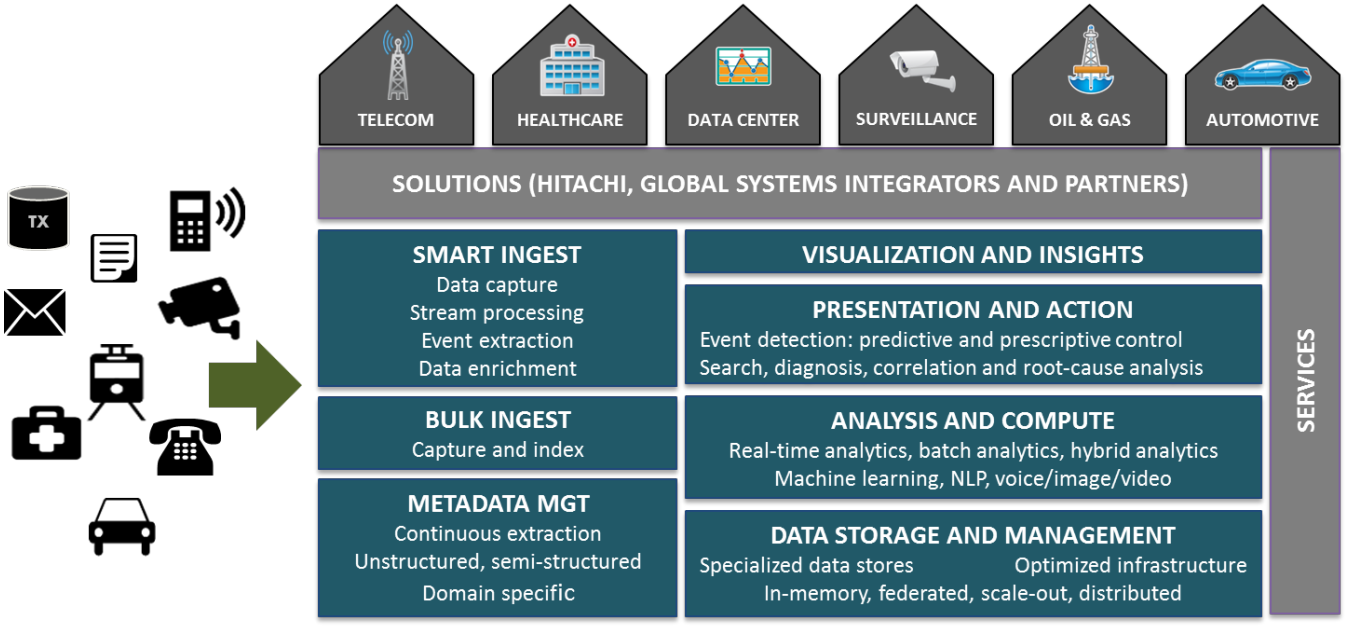

Services include:

- **Data ingestion services.** This layer provides techniques for ingesting data from a variety of internal and external sources. These data sources may be either structured, semistructured or unstructured.
- **Analytics and compute services.** Hitachi leverages a mixture of open source and in-house developed IP to deliver the services necessary for both batch and real-time analytics and hybrid combinations of both. Within the company's big data labs they are developing a library of foundational components to speed the development of applications. This library includes machine learning algorithms, spatial-temporal event detection, complex event processing, event extraction from unstructured data, root cause analysis and other techniques derived from developed solutions.
- **Data Management and storage services.** There are 2 layers of services. Data management services comprise the ability to manage a variety of different data sources, including traditional relational databases, temporal and graph data stores, key value stores and NoSQL databases. To support the speed and variety of data requires that data be stored either in a traditional file system or a distributed file system, a memory grid, or an object data store. Hitachi delivers optimized storage for these different types of data.
- **Visualization services.** Data visualization services support complex analysis to query, predict and correlate results from real time and batch analytics. The visualization services support root cause analysis to determine where problems exist and how they can be addressed. Hitachi works closely with visualization and integration tool partners to deliver these services.

Hitachi Analytics Solution for the Telecommunications Industry

To address the problems in the telecommunications industry as described earlier, Hitachi has created a big data solution based on its framework (see Figure 2) leveraging both in-house innovation and collaboration with technology partners, such as probe and deep packet inspection (DPI) technology providers. To improve network operations, HDS aggregates the real-time network streaming data combined with historical and business data that is siloed in many different data sources. In order to transform data so that it can be predictive and prescriptive, HDS synthesizes real-time analytics with historical analytics from past performance.

The collection of probe data (network metrics) or packet information can take place at different locations. These probes can be placed at the edge or core of the network transporting the information (data plane), or in the network transporting information relative to how communication sessions are established and managed (control plane). Hitachi's big data solution can collect, analyze and correlate information that emanates from many different parts of the telecommunications networks and efficiently scale to address the massive volumes generated by large telecom networks.

Hitachi's big data solution can collect, analyze and correlate information that emanates from many different parts of the telecommunications networks and efficiently scale to address the massive volumes generated by large telecom networks.

HURWITZ & ASSOCIATES
Insight to Action

Figure 2. Hitachi Analytics Solution for Telecom Is Built on Its Big Data Framework

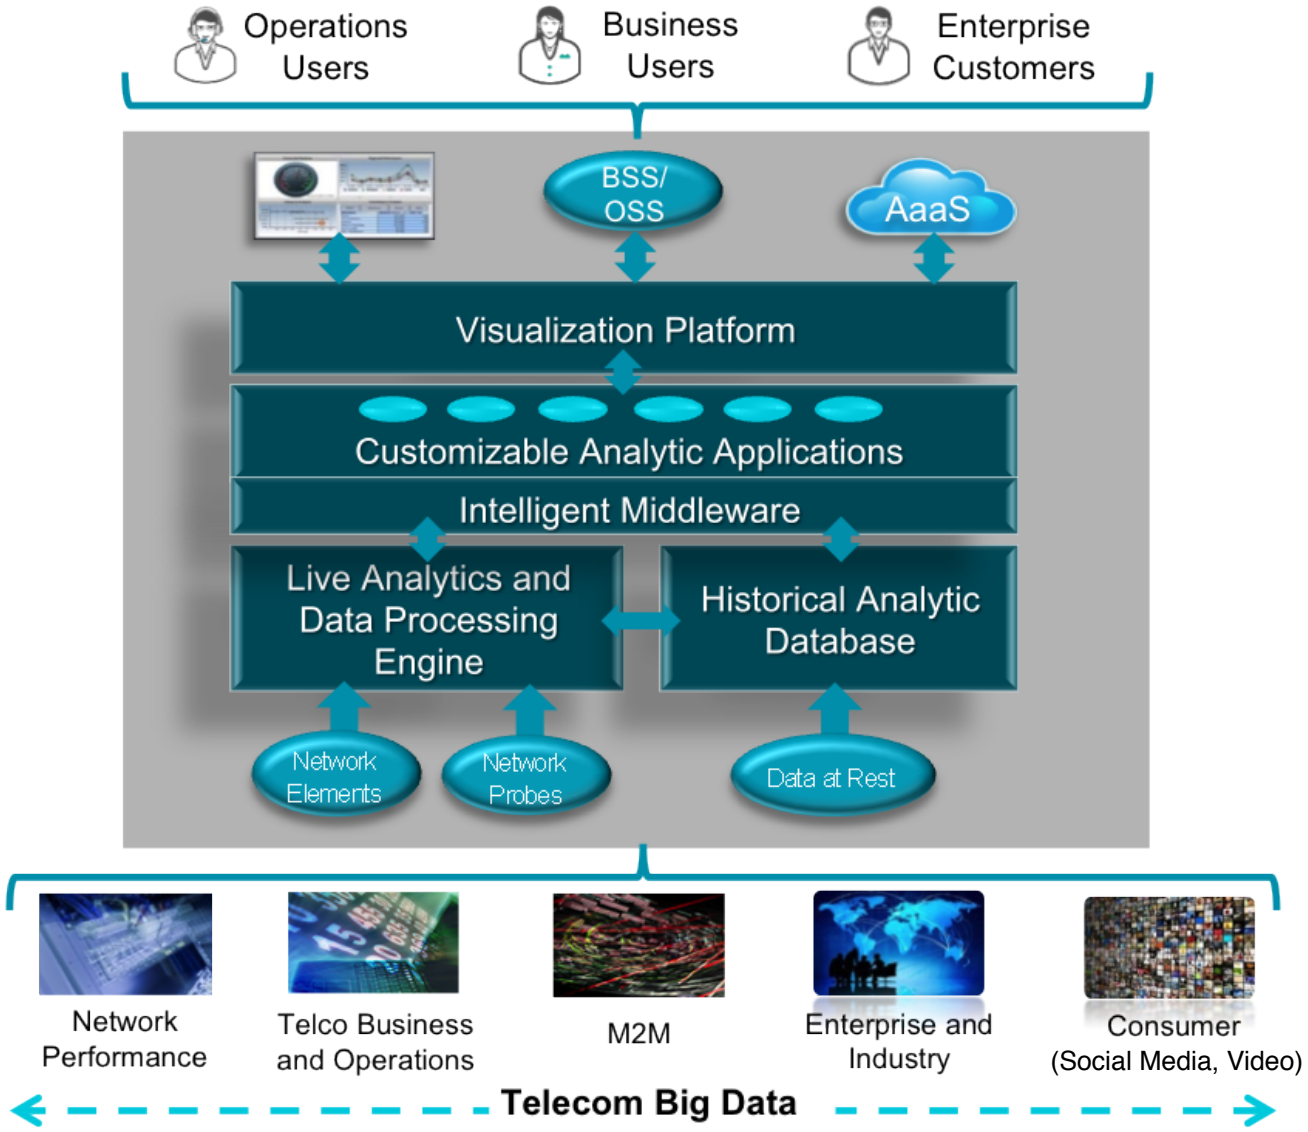

legend: M2M = machine to machine
BSS = business support system
OSS = operations support system
AaaS = analytics as a service

Leverage the Big Data Framework

Once the probes or network elements generate network traffic, the HDS live analytics and data processing engine is used to ingest this streaming data together with data from a historical analytical database so that the information can be analyzed. The framework provides the option of implementing either dynamic (real time) or opportunistic (batch) analysis. The result of this processing is placed into a visualization tool that can be viewed by analysts in

the dashboard. One of the benefits of the visualization environment is that it provides a way to see the current quality of service of the network in real time.

In order to provide the needed level of visibility and control over network probes, the stream-processing engine must capture and process the data through a series of adapters so that the data can be analyzed in real time. It then must capture the metadata and put the real-time data into the historical data store. The data store can then be used in conjunction with machine learning algorithms to detect the patterns of failures so that performance can be tuned.

Hitachi utilizes a unique persistent database store technology that is optimized for time series data. The data is stored in a graph database and can be efficiently searched and analyzed without needing to be decompressed first, resulting in ultra-fast query responses.

The telecommunications market is growing more complex with an increasing requirement to support many more use cases and more applications while maintaining the highest level of service.

The Value of Hybrid Analysis in Network Management

As we have pointed out, networks change dramatically over time and need to be able to support a variety of data types. Being able to collect this data and continue to learn from the experience of the network is critical to maintaining the right service and operational level in the network. By using machine learning algorithms, it is possible to begin to change the behavior of the network. For example, it is possible to correlate historical analysis based on time series to analyze how the network was behaving during a specific day, or over a month, to begin to identify what is normal and acceptable performance. This approach is preferable to looking at the average of a specific time period that will not take into account bursts in service levels that can impact overall quality of service.

Using this technique, it is possible to constantly monitor current data based on what is normal at that point in time. Therefore, it is possible to continue to tweak threshold rules proactively. If data analysis indicates that there is a problem, it is important to correlate network bandwidth with latency. The analytics engine identifies and correlates these relationships so that if performance is outside acceptable boundaries the operations team knows what actions need to be taken.

Conclusion

The telecommunications market is growing more complex with an increasing requirement to support many more use cases and more applications while maintaining the highest level of service. Customers of these network services are demanding quality and dependability of service and operations. Therefore, the ability to provide a framework and solution that can help organizations to have greater network visibility and control, plus anticipate and potentially avoid outages can be the difference between failure and success.

HURWITZ & ASSOCIATES
Insight to Action

About Hurwitz & Associates

Hurwitz & Associates is a strategy consulting, market research and analyst firm that focuses on how technology solutions solve real world customer problems. Hurwitz research concentrates on disruptive technologies, such as Big Data and Analytics, Security, Cloud Computing, Service Management, Information Management, Application Development and Deployment, and Collaborative Computing. Their experienced team merges deep technical and business expertise to deliver the actionable, strategic advice clients demand. Additional information on Hurwitz & Associates can be found at www.hurwitz.com.

35 Highland Circle • Needham, MA 02494 • Tel: 617-597-1724
www.hurwitz.com