How Will NonStop Fit Into the Internet of Things?

Justin Simonds >> Master Technologist >> Americas   Dean Malone >> NonStop Architect >> Caleb Enterprises Ltd.

PART I – Foundational Premises

There has been quite a bit of excitement about the Internet of Things (IoT) AKA Machine to Machine (M2M), Industrial Internet of Things (IIoT), Web of Things (WoT) and I'm sure there will be more identifiers that will emerge. It is the idea of machines or sensors providing information of interest to people, processes and other technology. This will usher in the 3rd wave of the World Wide Web.

The 1st wave involved the linking of information – documents to documents and search capabilities. This could be described as human to information. It was based on the ability to do quick research. Web 2.0 added many more human-to-human interactions with such web-based programs as YouTube, Skype, Twitter, Facebook, LinkedIn and the list keeps growing; the underlying element is person-to-person or person-to-group interaction.

This new and growing market under Web 3.0 uses the Internet as a large World Wide Service bus for passing machine information to interested end-points; be that other machines, humans or services/applications. How much information are we speaking about? IDC is forecasting the Digital Universe to be 44 Zetabytes by 2020 (Tera, Peta, Exa, Zeta – that's 10 followed by 21 zeros times 44). Consultants seem to be agreeing that 40% of that Universe will be M2M in 2020. Yikes! Even if they're off by a large factor, yikes!

Historically, the capabilities we are talking about for Web 3.0 have been monitored on a smaller scale by SCADA (Supervisory Control and Data Acquisition) systems that fundamentally monitor two types of telemetry points – control and analog. Control points are basically binary – on/off, open/closed, yes/no etc. Analog points typically have integer or floating-point values and set-points against them such as low-critical, low-warning, normal, high-warning and high-critical. Tied to those telemetry points are rules that react to them. Gartner has been talking of the integration of IT with OT (Operational Technology) like the SCADA systems just discussed. Businesses that have internal systems such as these - and most do - might consider integrating them as a first step toward IoT integration. With IoT we're talking about making expanded telemetry data available on the web and aggregating sets of them against particular entities such as:

- A member of a hierarchy such as a department, within a group within an organization.
- A person
- A geographical address
- A GPS coordinate

HP has its own particular vision of global SCADA. Please visit http://www8.hp.com/us/en/hp-information/environment/cense.html#.U_uOB_ldWSo for a brief overview of HP Labs CeNSE project (Central Nervous System for the Earth). The vision is to distribute nanoscale sensors around the globe to collect all sorts of information. To take an excerpt from the site:

> "By providing real-time information on the physical environment, the networks are intended to improve the way governments, businesses, and society respond to and manage environmental, biological, and physical/structural changes. Examples of potential CeNSE uses include roads, buildings, bridges, and other infrastructure; machines such as those used in airplanes and manufacturing plants; and organizations that work on health and safety issues, such as the contamination of food and water, disease control, and patient monitoring."

These are all inspiring ways of using the new technology and we are already seeing the rollout; one example being connected cars. OnStar generates monthly reports on covered cars: engine overview, transmission, oil life and even the tire pressure. This data is collected, transmitted and prepared as a report for its customers and this is just the beginning. Google already has a self-driving car with over 300,000 miles on it. Cars will communicate with each other; maintaining safe distances and drastically reducing - if not eliminating – accidents but the beginnings are already here.

Most people are aware of smart meters; the idea of providing feedback between water and power use within a home and trying to distribute requirements within a community. Let's not all wash our clothes at the same time, for example. Smart metering causes a great increase in volume and velocity. Back in the day (not so long ago) when a meter reader used to come around every month, a meter reading would generate 12 data points per year. To make things easy, let's assume an energy company has a million customers. That's 12 million readings a year to collect, generate bills for, report on and derive statistics and trends against. With smart meters installed, the company can now get a reading every 10 minutes. That's 6 per hour, 24 hours a day and 365 days a year; that is 52,560 readings per year. This is a three-orders-of-magnitude increase from the old 12 readings per year! With the same million customers the company now has 52.6 billion readings per year. That's a lot more volume, but it's also quite a bit more velocity because it is expected to be delivered in near-real time. What if the utility company decides it wants to read meters every minute? Every second?

Here are a few examples of how the information can be used:

- For each time interval, the consumption will be compared to the average mean of consumption to determine trend usage by customer, sub-station, grid sector and network.
- Customers who are consuming inordinate amounts of power that are outside the mean can be investigated for opportunities to conserve or move their consumption to a time when the network is under less load – particularly during peak power usage.
- Incentives can be provided to those particular customers who may have inefficient energy consumption appliances.
- Customers can be billed a higher rate for peak usage if it exceeds the norm by a particular percentage.
- Unusual occurrences can be flagged for exceptions. For example, a building's power going to ground or the sensor going offline might be indicative of the house being on fire so the fire department could be dispatched to investigate.
- Not all of this information needs to be stored in a database; just the data that is outside the normal range. Since the vast majority of the data will fall into the norm, this can radically reduce the amount of secondary storage needed to hold the trending data.

What makes this so exciting from a NonStop perspective is that most of this generated machine information will be in a very structured block of information. Sensors will not be generating snide comments, will not employ sarcasm, and won't use emoticons or any of the many things that makes textual, unstructured human information so hard to process for machines. It will be reliable. It will be structured. The length will be fixed. It will come in fast - possibly randomly and in bursts - and it will in many instances look very much like OLTP transactions, for which the NonStop system was designed.

At a recent talk Dr. Michael Stonebraker - a professor from MIT and a co-founder of Vertica, VoltDB and Postgres among others - said that 75% of all data stream applications will be OLTP-like. The rest will need Complex Event Processing (CEP) models. The IoT will create many if not most of the new data stream applications Dr. Stonebraker was describing. The NonStop operating system has always favored an OLTP/message-switch type processing load so this burgeoning market hits NonStop in its sweet spot.

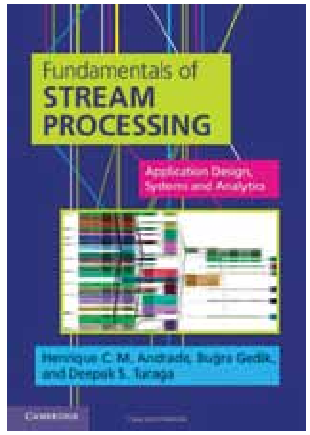

In the book "*Fundamentals of Stream Processing*" by Andrade, Gedik and Turaga, the first chapter contains a section titled "Towards Continuous Data Processing: The Requirements." We will be paraphrasing but encourage anyone interested to invest in the book. The chapter describes the challenges of ingesting, processing and analyzing information as it is continuously produced. The authors make the point that processing must keep up with ingestion rates; the velocity portion of Big Data or as we might state, be able to scale. Further on, they discuss parallel and distributed systems; clearly showing that a parallel architecture is ultimately the only way to achieve the scalability required.

Problems (or ingestion, processing and analysis) must be broken into sub-problems that can then be worked on and solved simultaneously. They must be broken into critical regions. The authors also state that the continuous nature of stream processing requires fault tolerance; although they also point out that different segments of stream processing applications may require different levels of reliability. As they phrased it, "For instance, tolerance to sporadic data loss is, in many cases, acceptable for certain parts of an application, as long as the amount of error can be bounded and the accuracy of the results can be properly assessed. Other application segments however cannot tolerate any failures as they may contain a critical persistent state that must survive even catastrophic failures." (ibid p. 6). The requirements for high velocity, critical segments of the stream processing environment are for a massively parallel, highly scalable, fault tolerant system. Sound like any system you know? From a requirements standpoint, NonStop starts looking like a pretty good contender.

The new processing requirements of IoT are largely concerned with velocity and volume. There will be a lot of data coming in very quickly. To use an overused analogy, it will be a fire hose – a data deluge. What architectural enhancements might allow NonStop to thrive in this new environment? InfiniBand is part of that new thinking. In Part II of this series we will discuss some of the technical underpinning which will enable NonStop to uniquely target the new IoT market.

Justin is a Master Technologist for the Americans Enterprise Solutions and Architecture group (ESA), a member of the HP IT Transformation SWAT team, and a member of the Mainframe Modernization SWAT team. His focus is on real-time, event-driven architectures, business intelligence for major accounts and business development. Most recently he has been involved with modernization efforts, Data Center management and a real-time hub/Data Warehouse system for advanced customer analytics. He is currently involved with HP Labs on several pilot projects. He is currently working on cloud initiatives and integration architectures for improving the reliability of cloud offerings. He has written articles and whitepapers for internal publication on adaptive enterprise, TCO/ROI, availability, business intelligence, and the Converged Infrastructure. He is a featured speaker at HP's Technology Forum and at HP's Executive Briefing Center. Justin joined HP in 1982 and has been in the IT industry over 34 years.

Dean is one of the pioneers of Message Oriented Middleware (MOM), having chaired three panels on MOM in 93, 94 and 95 at COMDEX. He developed the world's first fault-tolerant shared memory (XIPC on NonStop in 1995) deployed that product as the first customer implementation of active NonStop process pair (four programs implemented) and also ported Seer HPS/NetEssential 4GL-middleware to the NonStop. His biggest middleware achievement was the porting of IBM MQ-Series to NonStop as Chief Architect in 1998. He was the infrastructure architect for the Province of Ontario responsible for implementing the world's first wireless WAN-based mobile workstations for OPP, regional police, carrier enforcement and ambulance services. His customers include banks, brokerages, retail, EFT/POS switches, funds wire, vendor products, airlines, reservation systems, industrial automation and more. He has built systems on NonStop, VMS, Stratus, Unix and PDP-11 and has played roles as architect, technical lead and hands-on technical problem solver as a consultant for over 30 years. He is presently completing an RDMA Middleware product that will implement distributed shared memory, semaphores and queue-based messaging between NonStop, Linux and Windows servers over InfiniBand.