

## Powers of 10 – Size is only half of what matters

If you can measure it, we can historize it.

People often talk about size. “Mine’s bigger” is a frequent boast when people gather to have cocktails and talk about their new PC with 11 CPU cores, or their gargantuan movie collection (burned legally from their purchased DVDs of course), their 6-wheel drive Humvee 7 with Dual V-12 engines, or even their favorite Gillette razor with 10 blades (here, feel this).

Size does indeed matter. The Internet of Things (IoT) has ushered in an area without upper bound on the size of the ever-growing recoded history about the world around us. “Big Data” is now a convenient euphemism for this information explosion, and every technology provider is rushing to proclaim how their platform is “bigger” and “better” than their competition. To succeed in today’s Analytics-driven world, it is not enough to be bigger. You also have to be smaller!

Today, inmation makes the bold claim that we are smaller than all of them. Curious? Read on!

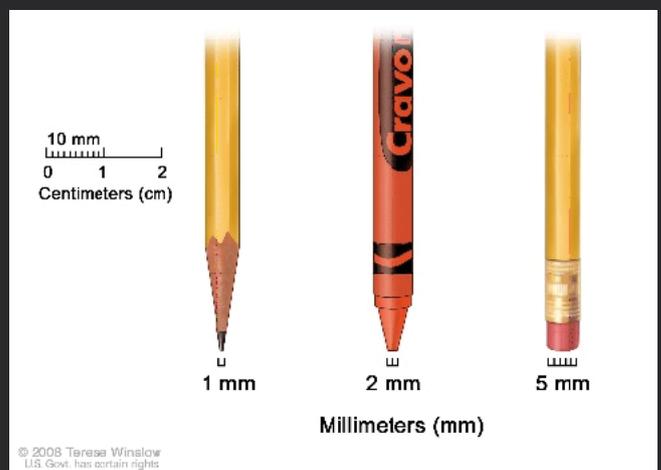
Most of the tech giants are busy beating their chests about the SIZE of the data they support, while quietly ignoring the equally important dimensions of SPEED and RESOLUTION. If you can store a lot of information, but have difficulty using it to solve real-world analytics problems, you have created a data graveyard, or “write-only memory” as one consultant so lovingly called a famous ERP vendor back end. Modern problems happen in real-time, and so big is great as long as it’s also FAST to retrieve and supports PRECISE time stamps.

inmation Historian v1.1 now has the SMALLEST time resolution of any commercial system on the planet. While we have always been upward scalable to the Petabyte or Yottabyte dimension in SIZE, we have now increased the internal timestamp precision to support PICOSECOND timestamps. Before v1.1, system:inmation was already a world-class historian, capable of storing data with a millisecond resolution. We now support any imaginable project requirement for data.

*“We can now deliver incredible historian capabilities to any industry or even any science lab on earth.”*

First things first. What’s a picosecond? Technically, it’s  $10^{-12}$  seconds, or a decimal point followed by 11 zeroes and then finally a 1. Light travels only about a third of a millimeter ( $1/3$  mm) in one picosecond. So, it’s a tiny little slice of time.

Who cares about picoseconds? Not much can happen in one, right? Au contraire.



A (huge) client approached us with a smart manufacturing opportunity involving discrete production. The machine lines are already capable of producing history at 1Khz (1000 samples per second per machine). Soon, the lines will produce history at 10KHz (10,000 samples per second per machine). With many machines per line and many lines per factory, millisecond resolution was not enough for this demanding application.

Nothing like a real-world challenge to get the nerdy juices flowing in our Server team. The results are spectacular! We have increased our typical historian data rate from 100,000 sustained inserts per second to over 1 million sustained inserts per second on commodity hardware (a regular off-the-shelf box in the 1000 dollar range). In addition, the data can be flexibly stored with milliseconds, microseconds, nanoseconds or picosecond resolution. We can now deliver incredible historian capabilities which SCALE UP to Yottabytes, SCALE OUT across any number of replica sets, STORE IN picoseconds and COLLECT in megahertz.

[system:inmation now exceeds requirements for every industry and every science lab on earth.](#)

## The Customer Details

The customer factory data has been de-identified for proprietary reasons, but the chart is real. This is a trend display of an eight-minute machine production cycle with 12 distinct manufacturing operations. For one operation, system:inmation collects a total of 1 million

individual measurements for 25 distinct items, which translates to approximately 40,000 measurements per item in this eight-minute machine cycle frame:

You can see any potential significance in this picture. A Statistical Process Control (SPC) software algorithm can use this data in real-time with advanced statistical methods to raise quality alerts or machine diagnostic interrupts.

You can drill in from the multi-operations chart above to analyze a single operation (or part of it) on a single machine. Users do not have to know about particular start and end times, as the system keeps track of the manufacturing context of each individual operation. Hence, the system is made for track and trace in industries where millions of products are manufactured in relatively short periods of time. Here is the trend chart displaying multiple readings for a three-second fraction of one machine operation. Further drill-down to the detailed sub-second (and even sub-millisecond) "world of machine life" is just a click away.



## Solution Scope

**Flexibility:** In addition to the native trending, system:inmation allows you to retrieve the data over multiple interfaces such as OPC HDA or JSON-based web services, in either raw or aggregated using 40 standard aggregates defined by the OPC Unified Architecture specification.

**Speed:** Does your process consist of hundreds of machines producing enormous datasets for each manufactured piece of product? No problem, system:inmation Historian v1.1 handles it.

**Data Integrity:** For in-process analytics it is absolutely important that data compression does not lose the information about a single raw measurement. We compress data not information using the MongoDB Wired Tiger engine.

**Extensibility:** Our in-memory Lua script lets you access and manipulate in-stream your inbound data using any real-time logic or mathematical algorithm. As such you may use our product as a perfect automation component with regards to SPC and APC.

**Reliability:** Through the distribution features of the product, you can roll out any advanced automation algorithm to worldwide production sites in no time.

## Call to Action

It's not enough to simply think about big data. It's time to embrace big, fast, precise and reliable data!

We invite all operators of large hadron colliders, fusion reactors, SETI radio telescopes or scientists working with hi-res seismic data or engineers responsible for any in-process high-speed manufacturing line to discover the new system:inmation.

[The most powerful IoT platform on earth.](#)

## About the author

Timo Klingenmeier

*Co-founder, GM and Technical Lead inmation*

Timo has a history of 25 years in industrial data processing. He started off in 1989 as a Software Developer writing first data historians for MS/DOS with such ancient tools like Turbo Pascal and dBase. From 1993 to 2003 he acted as general manager for the industrial system integration companies iDV and best solutions. After the acquisition by Matrikon,

Timo functioned as specialist for industrial software and head of German operations until 2010. In 2011 and 2012 he was consultant for international offshore wind projects. Since 2013 he focuses on the development of the NextGen information management software that is system:inmation.

