

open Historian

Version 2.0

Ritchie Carroll
Steven Chisholm

August 14, 2013



- C# openHistorian archiver is part of the openPDC
- Same “.d” data structure as TVA DatAWare
- Simple tools for data extraction
 - Data Extraction Tool
 - Data Trending Tool
- Real-time interface available via a web service

High performance – Over 200,000 points per second can be archived on a single server (at the lower end of server product lines)

High availability – Supports redundant and/or clustered deployment

Scalable – Clustered system deployment allows horizontal scaling to accommodate any volume of data

Proven – In production use for over a decade, with a single installation archiving over 5 billion points per day

Easy data extraction – Archived data can be exported in a variety of formats including COMTRADE

Compression available – Supports both lossless, full resolution and traditional sloped data compression storage

Automatic metadata synchronization – The openHistorian can be easily integrated with the centralized configuration systems

Massively parallel distributed computational support – Can be integrated with the Hadoop framework to allow fast data mining over many Petabytes of archived data

Lowest cost – The major cost for implementation of the openHistorian is for physical storage. There is no per-point or per-server license fees required.

Version 1.0 - Current State

- Very stable
- Millisecond time-resolution
- 32-bit floating point values (with quality)
- Performance limit around 100 PMUs
- 3x real-time replay speed
- Supports master/slave metadata modes
- Hadoop replication option
- New COMTRADE export option
- No standalone installation (i.e., deployed with openPDC)

Version 1.0 - Known Deployments

- TVA
- Entergy
- PG&E
- Dominion
- Every openPDC installation (via stats and/or active phasor archive)

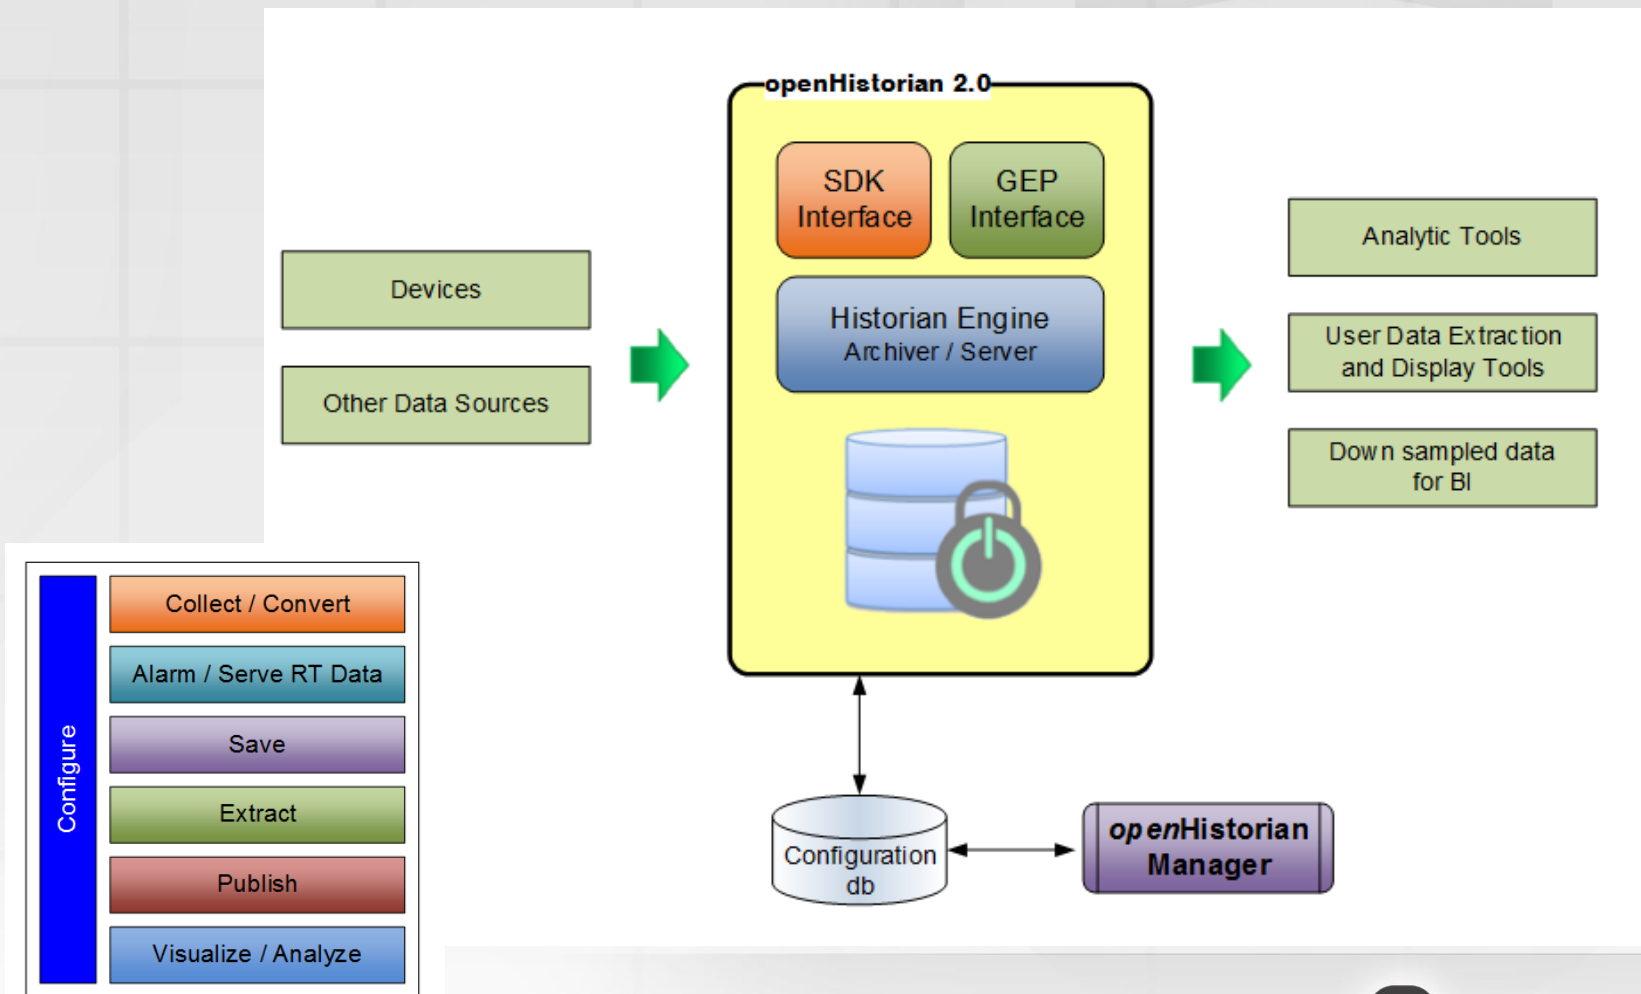
Why is Version 2.0 Needed?

- Performance improvements – especially read speeds
- Support for multiple data types while maintaining low storage requirements
- Standalone service based installation
- High-speed socket based historical data access API
- Real-time GEP based pub/sub data bus

Why does OG&E need Version 2.0?

- Speed !
 - High frame rate data refresh
 - Very quick query response times (60,000 points in less than 30 ms)
 - Quickly process very large data blocks e.g., post-process a day's data
- Ability to insert data out-of-sequence
- Overcome issues with current relational system

Version 2.0 - Architecture



Version 2.0 - Keys and Values

- The Key (3 – 64-bit unsigned integers)
 - Key 1: Point ID
 - Key 2: Time (Ticks: 100-Nanosecond Resolution)
 - Key 3.0:
 - First 32 bits -- Extended Time Precision
 - Second 32 bits – System Reserved Data -- e.g., leap seconds, duplicate entry counter
- Data (3 – 64-bit unsigned integers)
 - Key 1: Low Quad-word Data Value
 - Key 2: High Quad-word Data Value
 - Key 3: Quality

In practice, “128 bit data” means

- The measurement can be, for example:
 - 32 bit floating point
 - Unsigned 64 bit integer
 - 16 characters of text
 - Two, 32 bit floating point numbers (a complex number like a phasor)

... or any other structure that's needed

Version 2.0 - Current State

- Alpha version released
(and in pre-production at OG&E)
- Source code available from codeplex:
 - <http://openhistorian.codeplex.com>
- GEP and Socket API based integration operational

Version 2.0 - Next Steps

- Develop a robust high-level API that minimizes data loss during failover
- Expand available data types
- Add support for new protocols & interfaces – OPC, Modbus, etc.
- Develop rudimentary visualizations/tools
 - Trending tools
 - Process flow visualization builders
 - Data extraction / integration tools