EcoStruxure Geo SCADA Expert

(previously ClearSCADA)

Software for Telemetry and Remote SCADA Solutions

Performance Guidelines

November 2019

Schneider Electric

Introduction

EcoStruxure™ Geo SCADA Expert (subsequently referred to as 'Geo SCADA Expert') software is a high-performance SCADA system, capable of scaling to large databases and/or fast data processing. Each new release of Geo SCADA Expert software provides new features and higher performance to cope with the increasing demands of real-time monitoring, control, storage and analytics.
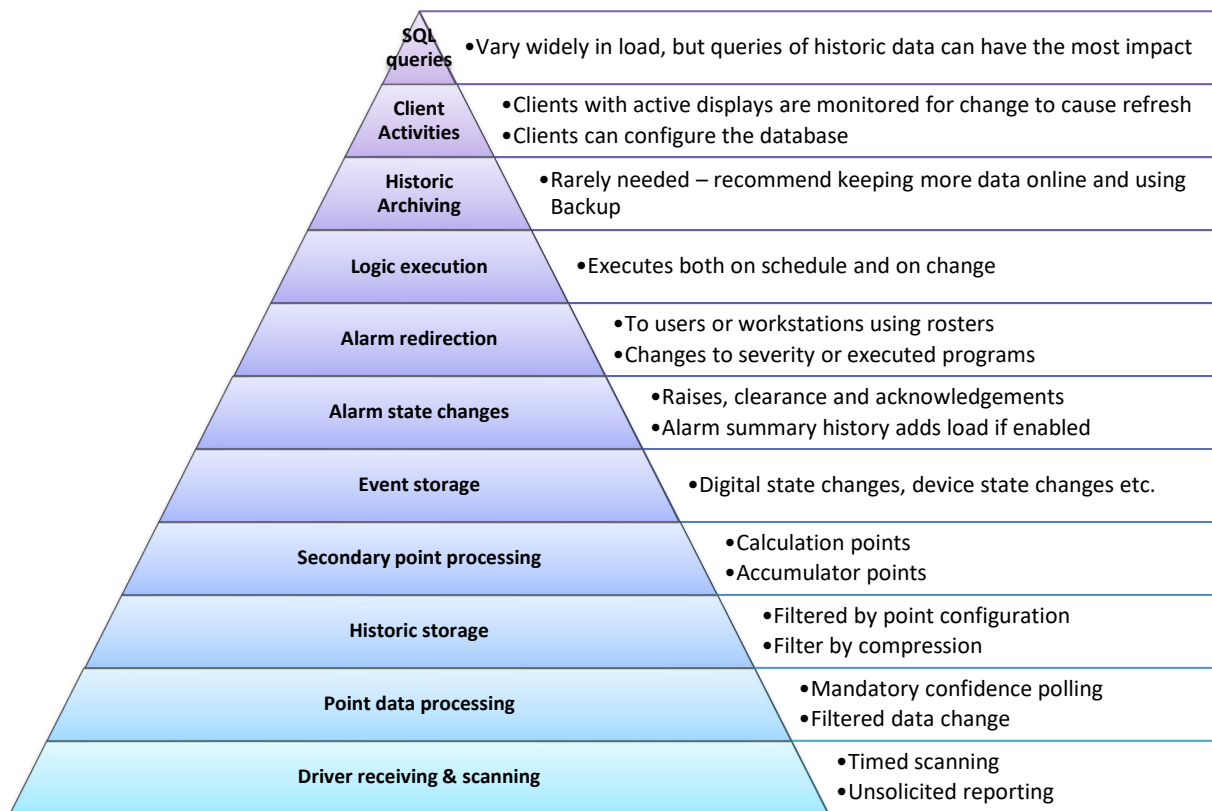
This document provides some guidelines for performance when creating and maintaining larger Geo SCADA Expert software deployments.

Geo SCADA Expert Software Server Activities

The ultimate limitations of a SCADA system are a combination of the physical resources of the system's computers - memory size and speed, disk access time and CPU speed - and of the database architecture which co-ordinates system activity. Having one fast physical resource such as disk storage is important, but the architecture will always be limited by the slowest resource. In high load conditions we find that the CPU and memory speed are limiting factors. This is typically the case with SCADA systems which are scanned and updated at high rates and is due to the need for a database to control updates using locking so that integrity can be assured.

Typical of most SCADA systems, the loading on Geo SCADA Expert software consists of many different activities, all of which are configurable to control the balance of work. A good way to visualize this is the pyramid below. The text to the right of the pyramid are notes to help give it context.

At the bottom of the pyramid is the basic receipt of data. This can be scanned using simple timed methods, or input by the devices using 'unsolicited' or 'asynchronous' data reporting. The latter methods are often better because there is no wasted activity reading unchanged data, although protocols such as Modbus^{TM} only have a scanning capability. As we move up the pyramid, the data processing quantity reduces, because each layer filters redundant or irrelevant data. For example, the driver may read all points in a given time, but the unchanged (or little changed) data is not passed to the next layer. Similarly, the historic storage layer will either filter by configuration of value dead-bands or by time intervals.

The pyramid diagram (from bottom to top):

Layer	Details
SQL queries	• Vary widely in load, but queries of historic data can have the most impact
Client Activities	• Clients with active displays are monitored for change to cause refresh • Clients can configure the database
Historic Archiving	• Rarely needed – recommend keeping more data online and using Backup
Logic execution	• Executes both on schedule and on change
Alarm redirection	• To users or workstations using rosters • Changes to severity or executed programs
Alarm state changes	• Raises, clearance and acknowledgements • Alarm summary history adds load if enabled
Event storage	• Digital state changes, device state changes etc.
Secondary point processing	• Calculation points • Accumulator points
Historic storage	• Filtered by point configuration • Filter by compression
Point data processing	• Mandatory confidence polling • Filtered data change
Driver receiving & scanning	• Timed scanning • Unsolicited reporting

As we go up the pyramid, while the number of processing steps decreases, there is often an increase in the amount of computational and storage workload. For example, the processing of historic storage requires more CPU cycles to store a new record in the database, and the processing of an alarm state change for a point will require new event log records and changes to many associated fields of a database record. The diagram is naturally not to scale, and there is significant variability in the workload of items at each level, but it is a useful guide to configuring a system. For example, historic storage data rates need to be controlled to avoid storing redundant data.

Alongside the activities listed in the pyramid, there are other activities the system runs in the background, such as the flushing of database information to disk and the synchronization of data from the Main server to the number of Standby servers in the architecture. This background load is approximately proportional to the aggregate of all the other activities in the pyramid.

Disk access time can restrict the performance of the system, and the activity types principally affected are:

- Queries for historic data - where a large amount of historic data needs to be searched and/or retrieved.
- The flushing of historic data to disk (including events and alarm summary records).
- Access (read or write) of large Data Table objects by SQL queries.

There is a specification of Geo SCADA Expert software limitations for some of these pyramid layers. This is detailed in the product help documentation under the heading and search term 'Operational Limitations - Server'.

This specification includes a limit to the number of historic data records stored per minute and per point, which is 4 - a minimum interval of 15 seconds. Such an interval does not prevent the field hardware from being read at a much faster rate. It could be 1 second for example. This is not an enforced limit, because Geo SCADA Expert software points can be configured to store data at this rate and higher, but there are follow-on impacts to configuring the system in this way.

While a typical application for Geo SCADA Expert software will need some data to change rapidly, that rapid rate is never required to be sustained. For example, analog trend data needs a high storage rate when the information content is high and a low storage rate when the information content is low.

Maximizing Performance

A question often asked is 'how many x can be processed in y seconds?' This is very difficult to answer, particularly because if that is the only activity, this number can be very high. But available capacity is reduced when other activities use system resources.

There are also some good practices concerning leaving headroom for unusually high activity if the system is exposed to unexpected conditions. It is important to design the system so that not only is there no data loss, but also that users of the system get an acceptable level of performance, such as display times and screen update times, particularly when the input changes are important to operators.

Geo SCADA Expert software is designed to avoid 'missing' any of the activities in the pyramid. For example, historic data is queued and cached at high burst rates, then stored later when there is time for the database to do this. Alarms and alarm redirections are also executed in a way to help them avoid being missed. Some parts of the pyramid can be configured to skip processing at busy times. Logic is the best example here. If a Logic program is queued from a state change, then that queue size can be limited to avoid unnecessary load. Also, some timed activities for scanning and Logic can be deferred if the previous timed activity is still being processed.

There are many variables which a configuration team can use to control system load. Measuring the CPU and disk activity of a server is not enough, due to the internal controls of database locking described above. While there can be no single number which guides all the activity of Geo SCADA Expert software a good measure which we have found effective is the total time the database spends in a locked state - in other words the time spent updating or reading data. It is important this is not too high so that other activities can take place and the database remains available for new activity to occur.

This total lock time can be found using the Geo SCADA Expert software Server Status tool. It is the sum of two numbers in the top line of the 'General | Locks' page. This line is labeled 'Database' on the left, and the columns '% Time in Excl Lock' and '% Time in Shared Lock' should be added. Note that this will be high on system start-up and should gradually decrease during continuous operation. It can be reset by a right-click menu. Total lock time can also be read by summing two OPC tags: '#LOCK.1.% Time In Excl Lock' and '#LOCK.1.% Time In Shared Lock'. You could create a Geo SCADA Expert software Calculation Analogue Point with the expression:

```
("#LOCK.1.% Time In Shared Lock" + "#LOCK.1.% Time In Excl Lock")*100
```

We recommend that this total lock time remains below 50% on the Main server for most of the duration of Geo SCADA Expert software operation. Naturally there may be busy periods, but the long-term average should be less than our suggested figure.

Using this lock time guideline may require a configuration team to try alternative strategies to get the most effective use of the system. In this case using the pyramid to understand loading will help to guide the team towards an effective balance.

The following table concludes this guide. It lists each element of the pyramid, how the volume of processing can be measured, and what steps could be taken to influence that amount.

Note that all items in the last column are advisory and can be exceeded, noting that the overall load is more important. The figures are based on large systems with high CPU, memory and disk performances. Your system may not achieve the figures listed in the table below, and each line's figures will be reduced by overall loading.

⚠ WARNING
SYSTEM PERFORMANCE DEGRADATION
Use of the diagnostic measurements described below and subsequent alterations to system settings and configuration can affect performance and data retention. We recommend testing measurements and changes on systems not connected to plant.
Failure to follow these instructions can result in death or serious injury.

Load Type	Description	How to Measure	How to Change	Estimated order of Magnitude
SQL queries	Vary widely in load, but queries of historic data can have the most impact	Query CPerformanceStats table. Select QueryExecuteTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. Server Status and Snapshot Log File - "Latest, Largest and Longest queries"	Ensure queries in Logic are marked NOCACHE unless necessary. Keep queries of history to short periods if possible.	Aim for fewer than 500 queries per minute, but fewer than 10 historic queries per minute.
Client Activities	Clients with active displays are monitored for change to cause refresh	Server Status and Snapshot Log File - OPC-DA Subscriptions	Review mimic complexity. Reduce or avoid indirect tag use (square bracket syntax). Reduce scripting if possible.	Design mimics with fewer than 500 tags (individual data items) in total, and avoid indirect tags, or use indirect tags only to objects which change seldom.
Client Activities	Clients can configure the database	Event log. Configuration change log (if enabled).	Configuration activity causes the configuration database to be processed and saved to disk.	
Client Activities	Logic can configure the database	See the column 'Write Config Count' in the System Query in ViewX 'Logic Execution Status'	Take care to reduce or eliminate Logic which writes to configuration fields of database objects. Not only does this cause data saves to disk and sync to Standbys, it also causes re-evaluation of validation.	Avoid using Logic which performs configuration regularly.
Historic Archiving	Rarely needed – recommend keeping more data online and using Backup	Find database objects of type 'Archive'. Check linked schedules.	Configure archive to work weekly, and ensure that it completes each week, to avoid a build-up of unarchived data. Consider not using the archive feature, and retaining history for longer on all servers.	
Logic execution	On schedule and On Change	See the columns 'Overruns', 'Execution Queue Max', 'Queued Executions Lost' and 'Execution Count' in the System Query in ViewX 'Logic Execution Status'	Look for overruns and reduce scheduled frequency. Look for queued executions lost and check for inputs which change too frequently and reduce their scan or data update rate. Look for high execution counts and review why. Spread execution times with offsets, particularly for multiple instances.	Aim for fewer than 200 logic executions per minute.

Load Type	Description	How to Measure	How to Change	Estimated order of Magnitude
Alarm redirection	To users or workstations using rosters. Changes to severity or executed programs	Query CPerformanceStats table. Select AlarmRedirTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. See server Status tool - page 'Database | Alarm Redirections'. Look for rapid changes in the list.	Review whether all states or severity levels need to cause redirection.	Aim for fewer than 20 alarm redirections per minute.
Alarm state changes	Raises, clearance and acknowledgements	Query CPerformanceStats table. Select AlarmRaisedTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. Review SQL query of alarm raise counts using: SELECT "FullName" AS "~FullName", "Id", "RaiseCount" FROM CDBPOINT ORDER BY "RaiseCount" DESC	Use persistence to reduce fleeting alarms. Review plant behavior to resolve plant issues.	Aim for fewer than 50 alarms per minute. Consider the burden on operator interactions. Also the total number of uncleared alarms can affect performance
Alarm state changes	Alarm summary history adds load if enabled	Check Server Configuration	Alarm Summary is a feature which is essential if you require Alarm Adviser. Alarm Summary can be enabled with the Server Configuration tool.	Enabling Alarm Summary on an existing system with high load will place it under further load - it will be necessary to review that load first.
Event storage	Digital state changes, device state changes etc.	Query CPerformanceStats table. Select EventsProcessedTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. To look at event counts from simple drivers, use "Select Id,FullName,PointEventCount from CPointSource". From ClearSCADA 2017 R2, see the column 'EventCountPreviousHour' of all objects.	For points which report excessive events, check and implement persistence and verify correct plant operation. If the event is not required for an object and/or state, reconfigure the Severity field to 'None'.	Aim for fewer than 500 events per minute. Also aim for fewer than 20 events per file granule per minute. The granule file stores events for a group of consecutively numbered objects, and the group size is set in the server configuration tool. Look to keep granule file sizes in line with limits in the documentation and the default granule size alarms.

Load Type	Description	How to Measure	How to Change	Estimated order of Magnitude
Secondary point processing	Calculation points	Query CPerformanceStats table. Select PointsCPntCalcTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. See and compare process counts using SQL: 'SELECT ID, FULLNAME, PROCESSCOUNT FROM CDBPOINT WHERE TYPEDESC LIKE 'Calculation %' ORDER BY "ProcessCount" DESC'	Review high frequency updates. Choose on-input or on-time processing appropriately.	Aim for fewer than 1000 calculation point updates per minute.
Secondary point processing	Accumulator points	Query CPerformanceStats table. Select AccProcTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. Look for processing overruns using 'SELECT ID, FULLNAME, OVERRUNCOUNT FROM CRESETTINGACCUMULATOR WHERE "OverrunCount" <> 0'.	Timed report intervals may be too short. Selecting 'Continuous' will cause processing for every report time. Accumulators are updated for each process of the input point. Check whether each update type (checkboxes) are really needed.	Aim for fewer than 1000 accumulator point updates per minute.
Historic storage	Historic values can be filtered by point configuration, and by compression.	Query CPerformanceStats table. Select HistWriteRawTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. To see individual point historic data quantities, use this query: 'SELECT P.ID, P.FULLNAME, H.UPDATECOUNT FROM CDBPOINT AS P JOIN CHISTORYBASE AS H USING (ID) ORDER BY "UpdateCount" DESC', and compare over a time period.	Review point configuration to reduce unnecessary data storage. Accumulators may need Reset values stored and not Report values. Points may need only significant change values and not report values stored. If it is necessary to update current values often, consider using historic compression, on the Historic tab of points.	Aim for fewer than 5000 historic storage items per minute.

Load Type	Description	How to Measure	How to Change	Estimated order of Magnitude
Point data processing	Point processing includes timed scanning of values, retrieval of logged data (as supported by the protocol) and confidence polling. Note that confidence polling forces historic storage, and is intended as a 'slow' backup scan, such as daily.	Query CPerformanceStats table. Select PointsProcessedTotal from CPerformanceStats. Execute this twice with a known time interval and observe the difference. To see individual point updates, use this query: 'SELECT ID, FULLNAME, PROCESSCOUNT FROM CDBPOINT ORDER BY "ProcessCount" DESC', and compare over a time period.	Review scanning times, reduce confidence poll intervals. Widen Significant Change, increase Persistence, and widen alarm Deadbands.	Aim for fewer than 10000 point data processes per minute.
Driver receiving & scanning	Timed scanning, logged data and unsolicited reporting.	Inspect driver log files and channel communications log files. These will indicate what is being requested from the device, and what data is received.	The driver is the first 'filter' for data updates, and the configuration should be designed so that unrequired changes are 'seen' by the driver but not passed up to the database. Review scanning times, reduce confidence poll intervals. Widen Significant Change, increase Persistence, and widen alarm Deadbands.	No specific limit - but consider splitting devices across multiple sets/channels to allow separate CPU threads to distribute load.

Conclusion

The Geo SCADA Expert software team is committed to providing a high-performance system. Each new release provides more features and more performance. Geo SCADA Expert software combines unique Telemetry/Remote SCADA features, best in class driver support, and a growing scalability to meet the challenges of the 'original' Internet of Things.

Performance tuning is an important part of both system design and maintenance. In addition to this document, you can also find the Geo SCADA Expert software Design Guidelines document, search the Geo SCADA Expert software online help and consult the Geo SCADA Expert software Resource Center at:

https://tprojects.schneider-electric.com/telemetry/display/CS/Home

Contact Information

Email: OrdersTRSS@se.com	**Phone (Direct Worldwide):** + 1 (613) 591-1943
URL: Geo SCADA Expert – Official Website	