Ubisense DIMENSION4™

Setting up DIMENSION4 logging at a production site

Wednesday, January 11, 2023

Contents

Setting up DIMENSION4 logging at a production site

For large scale applications in assembly plants, it is important to consider the size and bandwidth requirements for logging, and to have a plan for how to do analysis and audit.

The following describes how DIMENSION4 logging is structured, how to configure DIMENSION4 logging, how to retrieve logging data in DIMENSION4, how to size logging bandwidth and storage requirements for a DIMENSION4 system, and how to structure the system so that logging is done efficiently.

How DIMENSION4 logging is structured

Logging in DIMENSION4 has a cellular architecture. There will normally be a separate logging server for each location cell and another logging server for the site cell.

Logging data includes event (i.e. sensor measurement) data and tracing data. Logging data is all stored in a single file, in a binary format. This means that it is not possible to directly inspect 'trace log' files in the DIMENSION4 system.

How to configure logging in DIMENSION4

The **Configure logging** tab is the first tab in the LSC workflow. Logging can be done on any machine that is running a Ubisense local controller: for each of these machines, you must configure the pathname of the base logging directory and the maximum amount of data to be stored on the machine. Once a machine has had these properties specified, it is possible to set it as the logging server for the site cell.

Normally that will be all the configuration that is required. When location cells are created, they will simply inherit the logging machine used by the site cell. But for very large systems, or systems that operate across multiple remote sites, it is possible to assign separate logging machines to each location cell.

How to retrieve logging data in DIMENSION4

The **Review location events** and **View trace messages** tabs in LSC allow retrieval of trace messages. The same user interface is used to view data in real time or to view historical data.

Location events are only retrievable for individual location cells. There is no interface that supports retrieval of all events for the whole site. This is because the number and bandwidth of event data would be too great for normal uses. Trace data.is only generated at site cell level and so it is always retrieved at site cell level.

Location event data can also be retrieved using the **ubisense_ls_event_to_file** command-line program: this is how data are retrieved for offline performance analysis. Trace data are retrieved using **ubisense_ls_trace_to_stdout**.

How to size logging bandwidth and storage requirements for a DIMENSION4 system

A spreadsheet *LoggingDiskRequirements.xlsx* is provided which is designed to support a notional system where there is a mixture of data stored raw and compressed. It provides this view:

Sensor-to-server-to-log disk usage and bandwidth analysis

Total number of sensors	600 sensors
Number of vehicle tags visible	0 tags
Mean vehicle tag update rate	1 Hz
Mean measurements per vehicle beacon	5 measurements
Total number of tool tags	150 tags
Mean update rate	8 Hz
Mean measurements per tool beacon	5 measurements
Logging Period	86400 seconds (1 day)
Sensor-to-server-to-disk I/O throughput	**8 Mb/s**

Uncompressed logging data stored for each period

Total disk from vehicle locations	0 MB
Total disk from tool locations	75064 MB
Total disk from sensor per-ten-second info	5360 MB
Total disk from sensor per-minute info	721 MB
Total disk per period	**81 GB**
Proportion of log used by event data	**93 %**

Logging data storage requirements when compressed

Disk for compressed events	18766 MB
Disk for compressed trace	608 MB
Total disk per period (compressed)	**19 GB**
Disk requirement for 14 periods (compressed)	**271 GB**
Proportion of log used by event data	**97 %**

Enter the expected size parameters in the yellow boxes and the predicted resource requirements are in the rose boxes. This example is taken from a car plant where there are no car tags because linetracking PLCs are used (so there is no vehicle event data) and 150 tools are controlled. A total of about 8Mb/s of logging data will be generated; each day, 81GB of event and trace data is generated, 93% of which is event data; if we store 14 days of compressed data we will need an extra 271GB.

In principle we might run logging on separate servers for different sets of cells, in which case we would be able to fill in this spreadsheet for each individual server, by using the expected number of sensors and tags for the set of cells covered by that server.

How to structure a system so that logging is done effectively

From the previous example, we can see that the logging bandwidth and capacity requirements are quite significant and need to be planned for carefully.

It is important to note that logging is not mission-critical data, and in general it is bad practice to treat it as such. Thus it is bad practice to use expensive, high-availability stores to store logging data, and it is important to ensure that the logging data performance requirements do not impact the performance of the actual mission-critical data storage.

One useful technique is to use a separate server machine for storing logging data. If a controller is run on a suitable server then it can be configured for logging using the **Configure logging** tab; then all suitably-configured cells will be automatically logged by servers running on that server. This setup would guarantee that the performance overhead of logging data to disk cannot affect the mission-critical parts of the system.

The example below shows an overview of the structure of a system that uses a separate logging server machine. It shows the two servers, two client machines (one on the customer LAN, and one at some remote site), and the sensors. There are two kinds of data flow: real-time data is handled by a simple, unreliable UDP protocol; historic data (and controller configuration) is handled by the Ubisense protocols, which may be tunneled over a TCP connection.
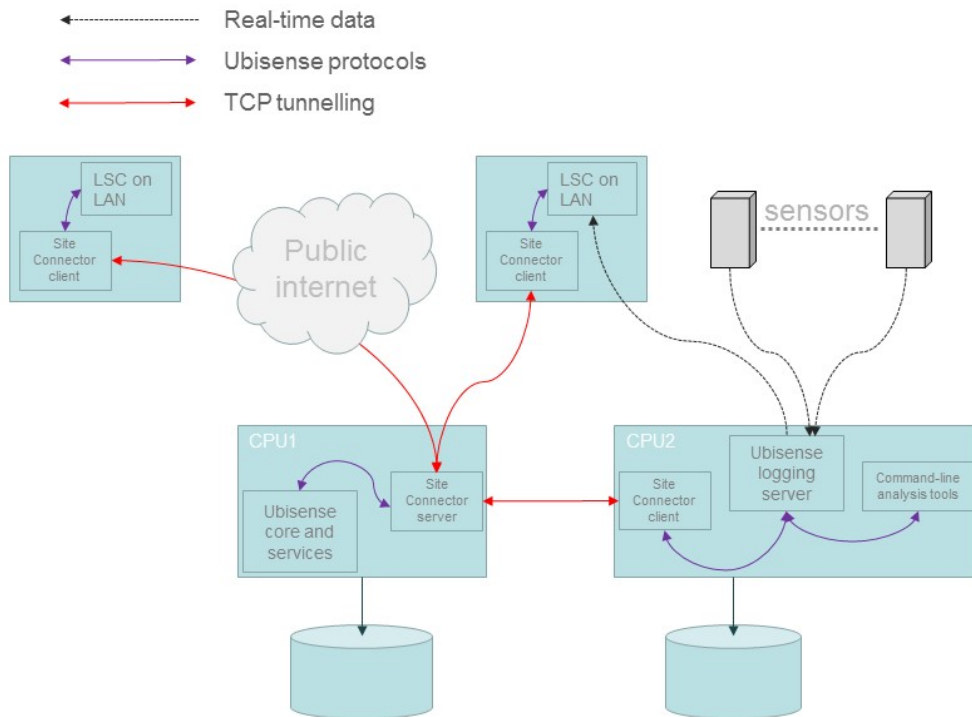
The important data paths to notice are the Ubisense protocols between the logging server and the analysis tools (run on the server) for historic analysis, and the real-time traffic between the logging server and the LSC on the LAN, for solver use and commissioning. Neither of these paths would touch the production server machine.

Frequently asked questions

Robustness of the system to logging server resource outages

Q: What happens when the logging directory is not available for a period?

A: Data gets into the buffer of the logging server. When the buffer is full, data is lost. Only UDP messages are sent from sensors to server, so there is no possible effect on the sensors.

Q: What happens when both logging directory and services are available, but there is disk latency? Will there be a queue building up? If so, where? Will this data be lost eventually?

A: Same as when the logging directory is not available. Yes, there will be a queue. It will be in the buffer of the logging server. Eventually this data will be lost when the buffer overflows. But again, only UDP messages are sent from sensors to server, so there is no possible effect on the sensors.

Q: What happens when the logging directory is available, but the logging services are down?

A: Data from sensors will be lost. Only UDP messages are sent from sensors to server, so there is no possible effect on the sensors.

Organization of logs on the disk

Q: Each line will be a separate Location Cell, so we will have many Location Cells. Does this mean we will have many logging services (at location cell level) writing into the same directory?

A: Yes. There will be 1 sub-directory per Location Cell.

Q: How many open files are there in a large DIMENSION4 system when logging?

A: In general the number of open files will be the number of Location Cells + 1.

Managing separate logging servers

Q: If using a separate logging server, would services have to be migrated to the second controller using Service Manager?

A: No, services would not be moved by Service Manager. If logging is configured to run on a separate server via LSC, the DIMENSION4 configuration services would ensure that the logging services are run on the correct machine.

Log retrieval from remote sites

Q: Can I use LSC (or **ubisense_ls_event_to_file**) to download historical logging data on a machine at a remote site over the internet, and what path does the data take?

A: Yes you can. The logging data will be downloaded using a remote procedure call which will be tunneled through the internet by Site Connector. Assuming the platform is in 'no multicast' mode, because the Site Connector server must run on the same machine as the core server, the traffic will always pass through the machine running the core server.

Q: Can I use LSC to access real-time logging data on a machine at a remote site over the internet?

A: No you can't. The real-time data stream is not available via Site Connector.

Log retrieval on a network

Q: Can I use LSC (or **ubisense_ls_event_to_file**) to download historical logging data on a machine on a network, and what path does the data take?

A: Yes you can. Assuming the platform is in 'no multicast' mode, the client machine will be in standalone mode, so the logging data will then be downloaded using a remote procedure call which will be tunneled through the network by Site Connector. Assuming the platform is in 'no multicast' mode, because the Site Connector server must run on the same machine as the core server, the traffic will always pass through the machine running the core server .

Q: Can I use LSC to access real-time logging data on a machine on a network, and what path does the data take?

A: Yes you can. The data will be sent directly from the logging server to the LSC process, via (unreliable) UDP messages. If the logging server is not running on the same machine as the core server then the data will not touch the machine that is running the core server.

Q: Can I use **ubisense_ls_event_to_file** to access historical logging data on the logging server machine itself, and what path does the data take?

A: Yes you can. The logging data will be downloaded using a remote procedure call which will directly access the logging service on the logging server machine. The data will not touch the machine that is running the core server.

Log compression

Q: Is there any best practice or tool support for compressing old log data, uncompressing it and loading it into a local system (apart from moving the data manually to a folder with exactly the right name etc.)?

A: There is currently no additional tool support.