Tiered Logging Requirements

Smart Grid

John Reynolds, CEO, Agile Fractal Grid

10-01-2020

What is the difference between a simple logging system and a Big Data “data lake” used for near real-time responsiveness of mission critical systems? They both record an ocean of data spewing out of sensors, but the need for random access to the data is a different problem. If the theater of operations is so vast that all the telecommunications capacity in the world today is not adequate to bring the trillions of event notifications to a single central point for analysis, what does one do? The answer is: a tiered collection mechanism that operates more like a distributed data lake than just a simple logging function. This white paper presents an approach that is being considered for the evolution of the smart grid around the world in the near future.

As has been described in other videos and whitepapers[1], the DDS protocol is very important for real-time processing for the electric power industry and is the new industry standard for the Open Field Message Bus for microgrids and substation interconnection where real-time control is now needed. The new architecture for delegating certain key decisions to the edges provides for a four-Tiered “laminar control” approach whereby rapid decisions for protection purposes are all made close to where the action is, but a “shared Consciousness” of what is going on is summarized and floated upwards. In this way, important matters are available for consideration at a national level, but that no one is overwhelmed by the volume of events taking place at 120 times per second at every end point.

[1] DDS Whitepaper provided by RTI

The DDS publish and subscribe protocol does this by configuring different domains for each Tier of the national infrastructure as is shown in the following diagram.

DDS Domain Map For Secure Management Channel

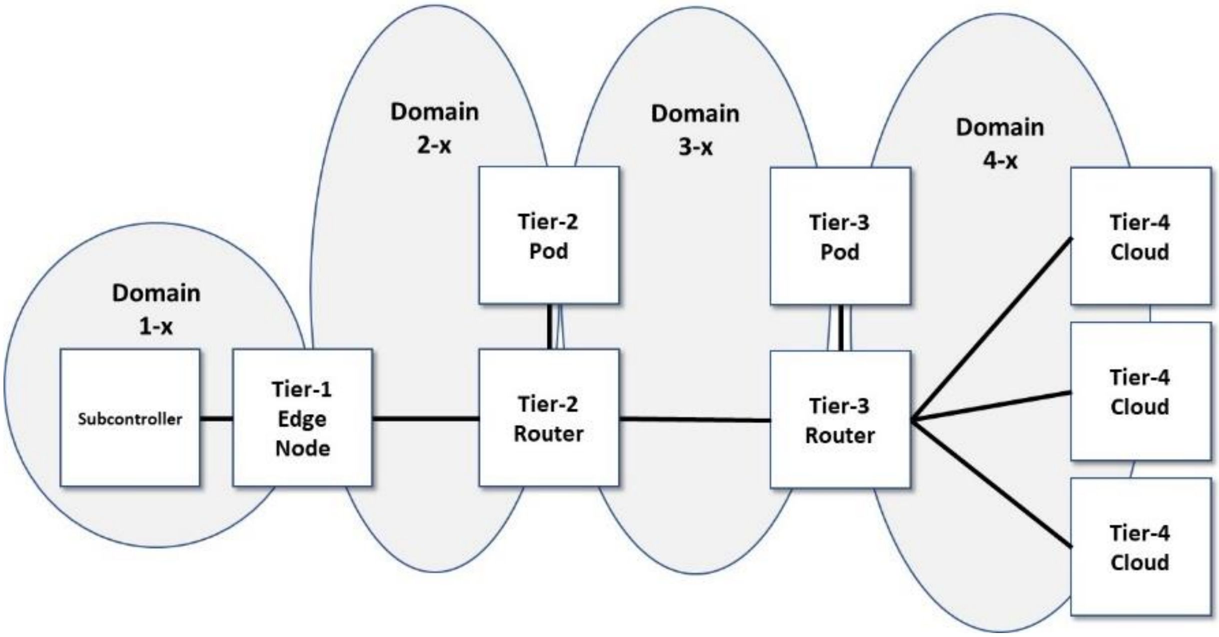

-1-

As events originate at electric grid subcontrollers at Tier-1 on the left of this approach, they are processed by recording their occurrence and applying rules for immediate action immediately at Tier-1 where time critical responsiveness is needed in about 4 milliseconds. That said, judgements are made as to how interesting the event might be to others. Critical events and summaries of others are then published on a Tier-2 domain bus and floated to the right.

At Tier-2, the events received are again logged and offered to listeners for rules and processes that are important for synchronizing the fleets of Tier-1 nodes that may be operating under the umbrella of the Tier-2 purview. These are normally operational in nature and are very useful when some catastrophe has occurred, and a Tier-1 node has malfunctioned, and the rest of the fleet must be adjusted to take up the slack. This is especially important to perform in just a couple of seconds, but the Tier-2 communicates “guidance recommendations” back down to the operational Tier-1 nodes that are still working. But again, the events reaching Tier-2 are logged, processed, and if they are deemed important enough for immediate analytical purposes, a new event is published on the Tier-3 domain and floated to the right so that they can be processed at a regional Tier-3 operations center.

At Tier-3, there are humans in the loop and a traditional situational awareness center can display faults, performance readings, and security intrusion information for human cognizance. Accounting and usage information is also collected here. The streaming logs of these events are treated as a part of the data lake available for analytical purposes. Severe events that are operational in nature tend to need to be recognized by a human operator with awareness in about 4 seconds. Events that affect financial positions in the energy markets for decisions relative to the 15-minute markets themselves tend to need resolution and decision withing 4 minutes. Such is the basis for the famous “rule of 4’s” in the energy industry.

For faults, measurement, and other information that needs to be shared on a national basis, the same forward technique can be used for publishing events and data for listeners at the Tier-4 level. By using this relay technique, the ability to listen and comprehend at Tier-4 is not overwhelmed by the billions of events that are occurring down at the Tier-1 level.

So, the tremendous volume of events is comfortably controlled by this relay system. But some forms of analysis up at the top may actually need to look back at some of the original low-level data for root cause analysis purposes. Thus, the logs that are kept at each of the different Tiers are more than just streaming logs. They need to be time sequenced databases, too.

The structure of the multiple tier capture needs to have the characteristics shown in the following diagram using the Tier- 1 Node as an example.

Tier-1 Edge Node

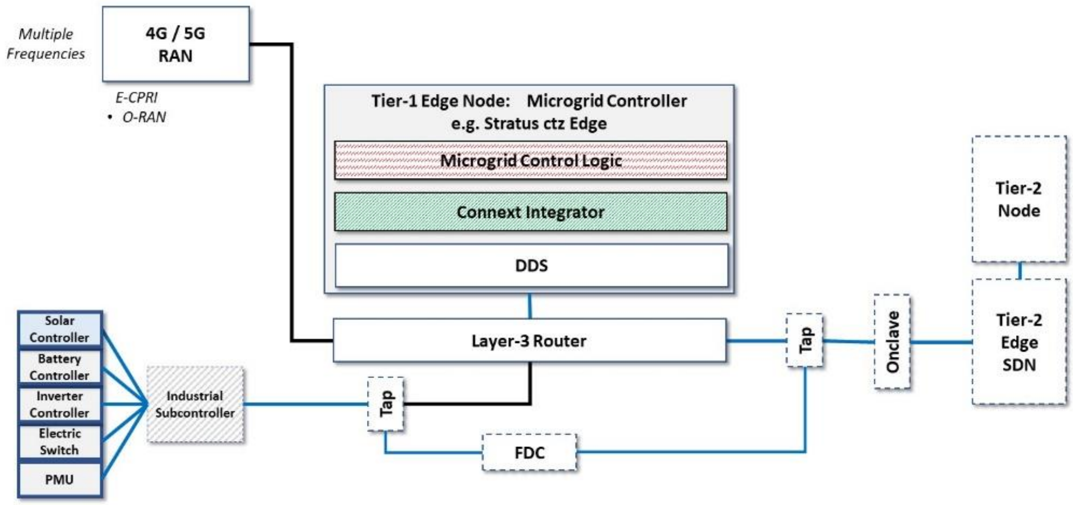

-2-

The interconnection from any industrial subcontroller needs to be delivered to the Tier-1 Node using the fault tolerant (multi-path) DDS protocol. In mission critical situations, it helps if the Tier-1 Node itself is fault tolerant such as with the Stratus ztc Edge Node. (In the diagram, the other communications and cybersecurity elements are also shown.)

To handle protocol normalization and data capture for logging purposes, a facility much like the RTI DDS Connext Integrator is useful at each Tier to serve both as the capture front end, but also the mechanism for distributing summary event northbound to the next Tier. The structure of the Connext Integrator is shown in the following figure.

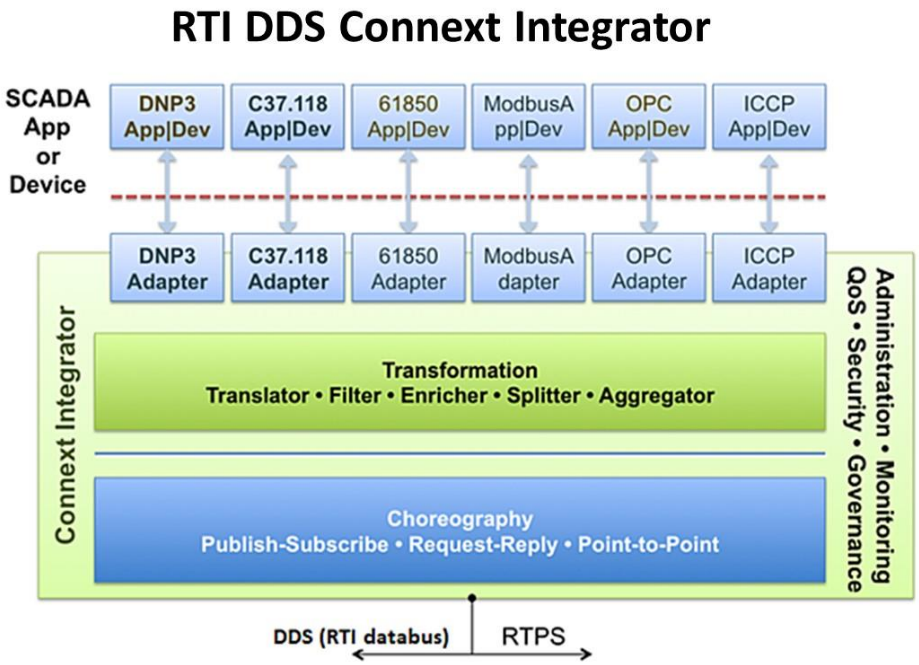

-3-

Within the Tier-2 Capture Processor, the “log” is actually a time series streaming data capture database organized for that the capture process both logs the original data that was received, and also that a validated copy of the input streams used to edit out anomalies in the data that would be harmful to processing. In addition to the capture function itself, the usual device management functions need to be in place as in any distributed systems node.

The detailed view of the Tier-2 Capture Processor is shown in the following diagram.

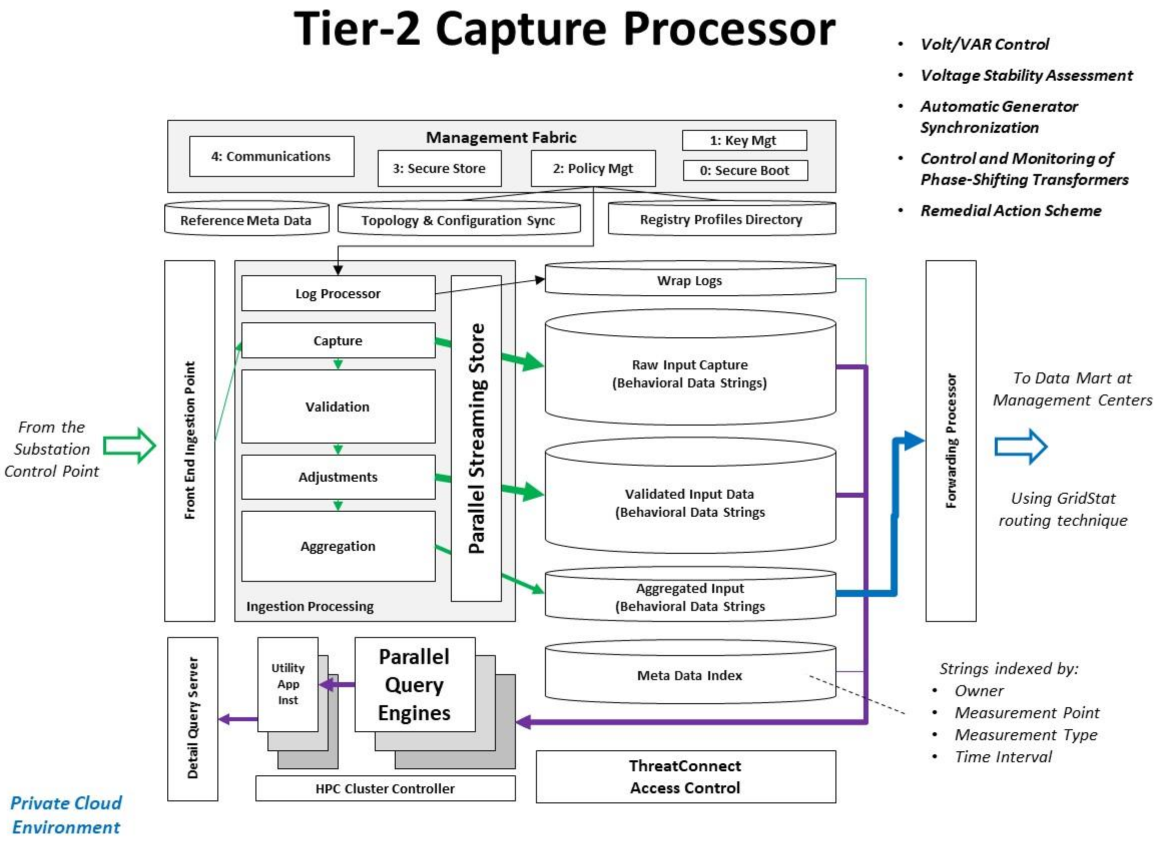

Tier-2 Capture Processor

- *Volt/VAR Control*
- *Voltage Stability Assessment*
- *Automatic Generator Synchronization*
- *Control and Monitoring of Phase-Shifting Transformers*
- *Remedial Action Scheme*

Management Fabric
4: Communications
3: Secure Store
2: Policy Mgt
1: Key Mgt
0: Secure Boot
Reference Meta Data
Topology & Configuration Sync
Registry Profiles Directory
From the Substation Control Point
Front End Ingestion Point
Log Processor
Capture
Validation
Adjustments
Aggregation
Ingestion Processing
Parallel Streaming Store
Wrap Logs
Raw Input Capture (Behavioral Data Strings)
Validated Input Data (Behavioral Data Strings
Aggregated Input (Behavioral Data Strings
Meta Data Index
Forwarding Processor
To Data Mart at Management Centers
Using GridStat routing technique
Detail Query Server
Utility App Inst
Parallel Query Engines
HPC Cluster Controller
ThreatConnect Access Control
Private Cloud Environment

Strings indexed by:

- *Owner*
- *Measurement Point*
- *Measurement Type*
- *Time Interval*

As we will see when looking at the larger picture, the data is not only captured here, but it also needs to be available for retrieval for intermediate results by inquiries initiated upstream, but assisted by parallel query engines resident right in the node for retrieval purposes.

The capture logs are arranged as parallel streams segregated by the source device being served. These string files contain blocks of data captured for storage efficiency, but the blocks are time sequenced with a multi-dimensional index that can retrieve blocks of events by the time that they were captured.

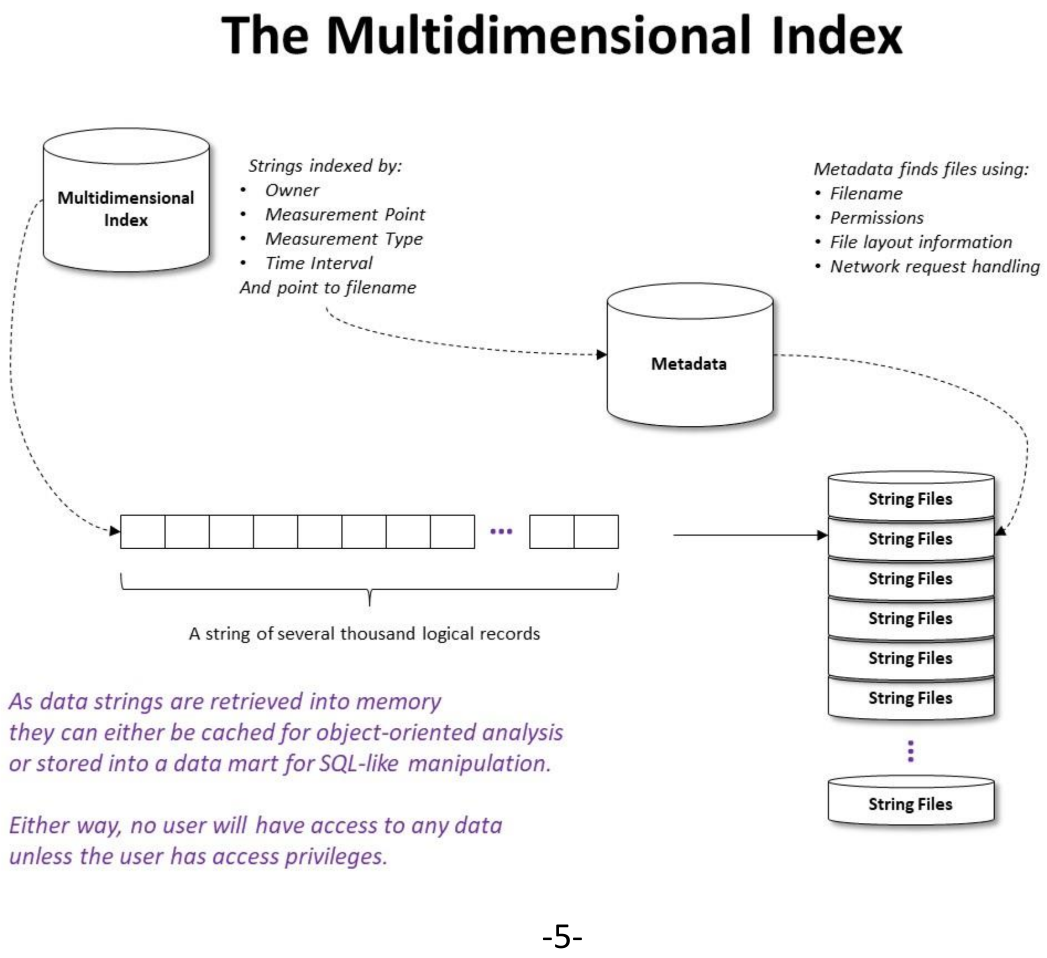

The Multidimensional Index

Multidimensional Index

Strings indexed by:
- Owner
- Measurement Point
- Measurement Type
- Time Interval

And point to filename

Metadata finds files using:
- Filename
- Permissions
- File layout information
- Network request handling

Metadata

String Files
String Files
String Files
String Files
String Files
String Files
String Files

A string of several thousand logical records

As data strings are retrieved into memory they can either be cached for object-oriented analysis or stored into a data mart for SQL-like manipulation.

Either way, no user will have access to any data unless the user has access privileges.

-5-

Typically, the device being monitored has a particular owner who has different levels of confidentiality for the data and different permissions are in play for reading, changing, or deleting this data, so the metadata associated with captured events is also a factor in the navigation of the captured data.

Thus, the data capture at its simplest if just a log function to some. But in the fullness of operations management, the event data captured, summarized, and forwarded for shared consciousness is really a distributed database.

The following diagram shows the relationship between the edge capture data storage at Tier-2 on the left, and the summary level regional capture of data at the Tier-3 Data Hub on the right.

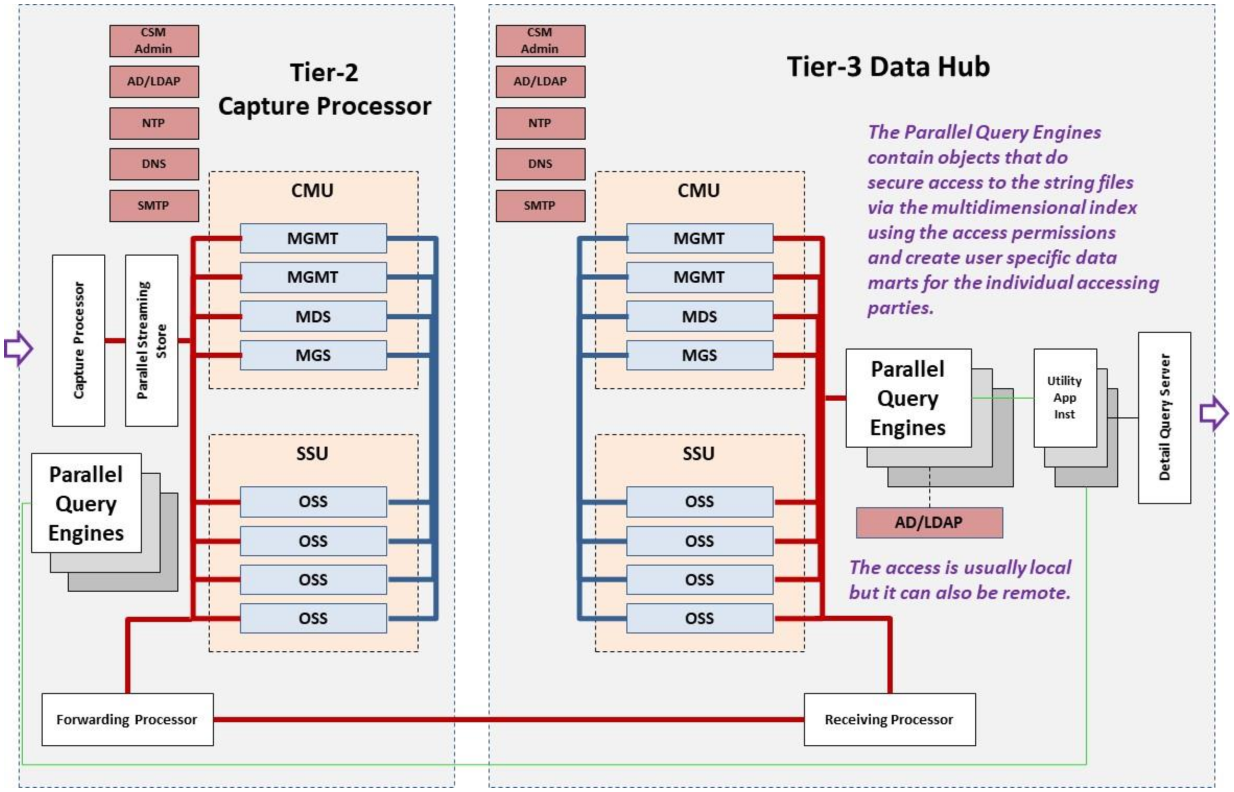

The Parallel Query Engines contain objects that do secure access to the string files via the multidimensional index using the access permissions and create user specific data marts for the individual accessing parties.

The access is usually local but it can also be remote.

-6-

At the top Tier, one wants to allow analysis of all the data available and permissioned to the analyst looking into longer time horizons. It is an ocean of data. And with the ability to run summary level investigations at the top very quickly. But when the real nitty gritty information is needed that is only available at the edges on the left, the parallel processing query should avoid trying to bring all of the data to the top-level query workstation, but instead partition the query

such that the power of all the processors at the edge do the first pass of the detailed database, and have the small amounts of the intermediate results set delivered back to the top Tier for final correlation, and then the workstation provide for the presentation. Using this massively parallel processing technique, an ocean of data can be queried and processed in a flash.