

Maximising performance in data-critical networked systems

By Dr. Rajive Joshi, Real-Time Innovations

This article shows ways how an implementation of the DDS specification can impact the three critical characteristics of data communication over a distributed network – reliability, predictability, and flexibility.

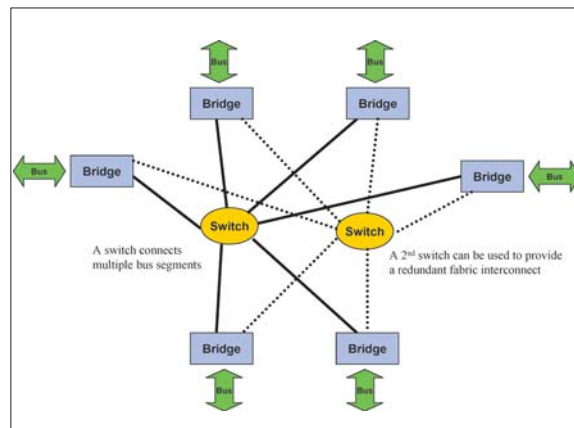


Figure 1. Switched fabric architecture. Multiple switches can be used to expand the fabric and provide hardware redundancy.

■ High performance and predictability are prerequisites for any large-scale networked system dependent on real-time data processing and analysis. Data representing actual events or system status must be evaluated while it is still relevant to tactical conditions, making it imperative to know when specific data is available, and to aggregate and evaluate that data in real time. Unreliable receipt times make effective analysis difficult or impossible.

Fast and predictable performance is always an issue in the design of any real-time embedded system, but especially so when designing distributed systems with thousands of nodes that need to move a lot of data around quickly in a dynamically changing environment. Switched-fabric networks can provide fast and highly scalable hardware solutions, and are now being increasingly used in such applications. What is needed beyond that is a software solution for bringing predictability, flexibility, and reliability to distributed data communications. Fortunately, a methodology known as publish-subscribe using a switched fabric network design is good at exactly this kind of problem.

Many large-scale data-critical applications can be characterised by three attributes - the need to gather and distribute data in real-time, the large amount of data being transferred, and the varied entities involved in this data exchange

that may change over time. For instance, air traffic control, financial transaction processing, battlefield, naval command and control, or industrial automation systems all are examples of data-critical systems which have these three attributes. These systems are not necessarily hard real-time, but their predictability requirements are an integral part of the functions they perform. They gather data from a variety of sources, sensors for example, and they distribute the data to a variety of users like databases, display devices or control algorithms. Furthermore, by their very nature, they are distributed.

Today's bus-based architectures, typically a multi-CPU VME backplane solution with hard-wired I/O interfaces to sensors and effectors, fall short in several areas in addressing the needs of data-critical systems. For example, these hardware transport mechanisms do not scale, are difficult to make fault-tolerant, and are difficult to modify and upgrade once they have been deployed. For these reasons, designers of complex, data-critical distributed systems are turning to switched fabrics to replace bus backplane and serial interconnect technologies. StarFabric, PCI Express Advanced Switching, Serial RapidIO and InfiniBand are some commercial products that implement different switched fabric designs. A switched-fabric bus is unique in that it allows all nodes on the bus

to logically inter-connect with all other nodes on the bus. Each node is physically connected to one or more switches. Switches may be connected to each other. This topology results in a redundant network or "fabric", in which there may be one or more redundant physical paths between any two nodes. A node may be logically connected to any other node via the switch(es). A logical path is temporary and can be reconfigured, or switched among the available physical connections. Switched fabric networks (figure 1) can be used to provide fault tolerance and scalability without unpredictable degradation of performance, among other features.

A key characteristic of switched fabrics is that they allow peer-to-peer communication between nodes without having to physically connect every node to every other node. With every node physically connected to every other node, adding a new node is exponentially more and more expensive as the number of nodes increases. Because a switched fabric network employs switching to achieve logical connectivity and reconfigurability, these systems can be architected to be highly scalable.

On the software side, publish-subscribe communication systems map very naturally onto switched fabrics. Publish-subscribe systems work by using endpoint nodes that communicate with each other by sending (publishing)

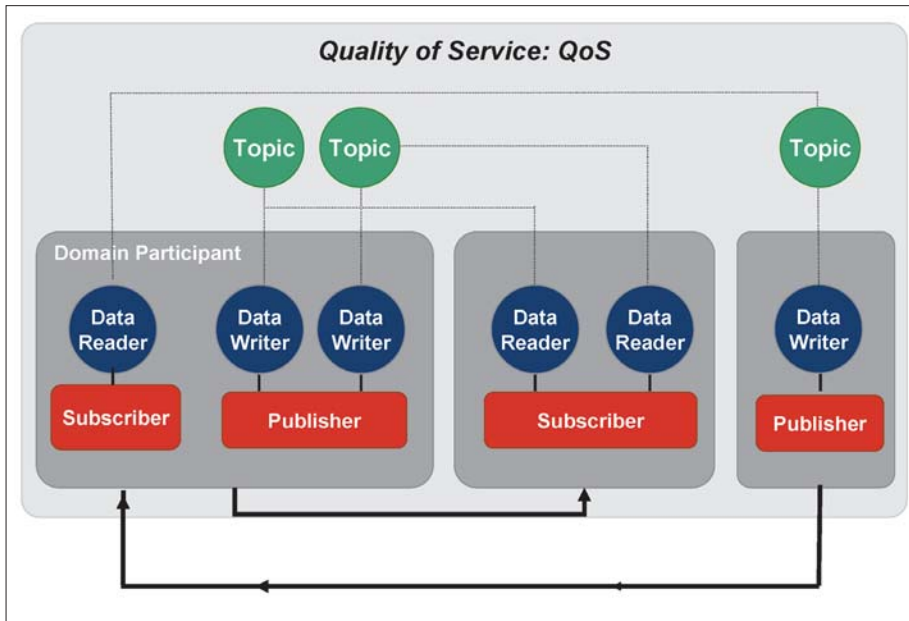


Figure 2. The Data Distribution Service data-centric publish-subscribe architecture organises the data in a distributed system around topics.

data and receiving (subscribing) data anonymously via topics. A topic is identified by a name and a data type. A data producer declares the intent to publish data on a topic; a data consumer registers its interest in receiving data

published on a topic. The middleware acts as the glue between the producers and the consumers; it delivers the data published on a topic by a producer to the consumers subscribing to that topic. There can be as many topics as need-

ed, a producer can publish on multiple topics, and a consumer can subscribe to multiple topics. The middleware layer isolates the data producers from the consumers; they do not require knowledge of each other.

A publish-subscribe software architecture allows producers and consumers to be loosely coupled. As a result it is naturally scalable, and can easily adapt to the changing needs of distributed data-critical systems. The producers and consumers are peers - they directly communicate with each other, so the topology of publish-subscribe systems can be closely matched to that of switched fabric systems. Thus, a publish-subscribe middleware layer can fully exploit the potential switched fabric network hardware.

The Data Distribution Service (DDS) standard (see box) specifies a data-centric publish-subscribe middleware layer, developed with the needs of distributed data-critical applications in mind. A well-designed DDS middleware implementation can be good at real-time data distribution, be easily field-upgradeable, and be transport agnostic. It can be better at real-time data distribution because publish-subscribe is more efficient than the traditional request/reply

The DDS Standard

The Data Distribution Service (DDS) for Real-time Systems standard, from the Object Management Group (OMG), defines a publish-subscribe system that has high performance, is efficient, and offers a predictable way of meeting the data distribution requirements of data-critical systems with minimal overhead. The standard can be found on the OMG web site at http://www.omg.org/technology/documents/formal/data_distribution.htm.

The DDS standard has been in existence since 2003, is actively maintained (currently at version 1.2), and maps very naturally to the topologies and capabilities of switched fabrics, and provides a solid technical approach to managing data distribution across large-scale distributed networks. The DDS standard has three main goals:

- To define a model for communication as pure data-centric exchanges, where applications publish (supply or stream) data which is then available to remote applications that are interested in it.
- To provide a mechanism for specifying the available resources and providing policies that allow the middleware to align the resources to the most critical requirements, giving system designers the ability to control Quality of Service (QoS) properties that affect predictability, overhead, and resource utilisation.
- To permit systems to scale to hundreds or thousands of publishers and subscribers in a robust manner.

based architectures in both latency and bandwidth for periodic data exchange. Further, it can be easier to upgrade in the field because publishers and subscribers do not need to care who or how many their counterparts are. And finally, since the middleware is layered on top of the physical means of getting the data from one place to another, it does not need to depend on the network transport or topology used.

Figure 2 illustrates the DDS data-centric publish-subscribe architecture. A topic has a name and a data type associated with it, and represents the application data model. DataReaders and DataWriters are associated with topics. A DataWriter can publish data on its associated topic; a DataReader can subscribe to data on its associated topic. DDS middleware automatically and anonymously sets up direct data flows between DataWriters and DataReaders associated with a topic, resulting in scalable and fault-tolerant data distribution. The DDS middleware quality of service (QoS) provides a means to tailor nearly every aspect of the behavior of the data-flow to the application requirements.

This marriage of switched fabrics and DDS real-time middleware offers architects new flexibility in adding capabilities that were once much more difficult to achieve. Many of the features offered by switched fabrics have complementary capabilities in the DDS-compliant middleware. For example, switched fabrics typically offer rich error management features such as the ability to recognize, report and route around failed paths. With DDS-compliant software, system designers can also take advantage of DDS error reporting facilities. A key feature of switched fabrics is support for multiple paths between nodes. This gives system architects the

ability to easily implement multiple physical interconnects that can be combined with sophisticated error management. Likewise, with DDS, applications can take advantage of redundant publishers that have different strengths or bandwidths. When a higher strength publisher fails, a lower strength one is automatically switched in by the DDS middleware. In addition to fault tolerance, this can also help with load balancing on heavily-used networks.

Switched fabrics specifications already provide for a hot plug or hot swap capability. This hardware capability can be combined with a “virtual” hot plug capability at the application level using DDS middleware. Unlike traditional tightly coupled client/server architectures, DDS middleware allows producers and consumers to be dynamically added or removed in an operational system. Many switched fabrics provide sophisticated features that allow, for instance, bandwidth-reserved, isochronous transactions across the fabric, something that is not supported by, say, Ethernet. Corresponding to the hardware QoS (Quality of Service) facilities, DDS-compliant middleware can offer a number of QoS policies that make predictability at the application level possible. For instance, the `TRANSPORT_PRIORITY` policy allows developers to manage how they prioritise one data flow over another.

The existence of DDS as a standard specification makes it ideal as part of a COTS-based solution. However, the value of the solution is highly dependent upon its implementation. The DDS specification defines certain features and capabilities, but not how they should be implemented. A carefully designed middleware architecture can reduce the likelihood of a fault, limit the damage of a fault if it does occur, help

detect faults immediately, protect the middleware from errors in application code, and isolate applications from errors in other applications.

That architecture can also deliver significant advantages in the performance and flexibility of network-distributed data communications. For example, the DDS specification defines how a publish-subscribe communication model should work for a distributed real-time network. The DDS specification defines DataWriters for publishing and DataReaders for subscribing to a single topic on a user-defined data type. This in itself is standard and straightforward, but how this is implemented can have a significant impact on network performance and scalability. A robust implementation improves both performance and scalability by defining an architecture that supplies each DataWriter or DataReader with a queue that buffers messages bound for another endpoint through a transport. This architecture supports direct end-to-end messaging, since each endpoint (a DataReader or DataWriter) in each application directly communicates with a sister set of endpoints. Each endpoint has a dedicated set of buffers to hold messages in transit to other endpoints. This queuing architecture provides for an optimised transfer of messages from DataWriter to DataReader, no matter where each resides on the network. And because the endpoints queue and buffer transmissions to other endpoints, this architecture can easily scale to large and complex networks, still with predictable delivery times.

In a similar manner, DDS defines the concept of a DomainParticipant, which is the fundamental container entity that can participate in a publish-subscribe network. A DomainParticipant can contain many DataReaders and DataWriters. Typical applications may use only one domain, and therefore have one DomainParticipant. However, applications are free to create several DomainParticipants so multiple instances of this entity can exist simultaneously. Multiple execution threads are a way to optimise responsiveness and performance, while also allowing the system to scale across a broad fabric-based network. One possible approach is to use several dedicated threads for each DomainParticipant, in this manner:

Event thread – DDS allows application designers to associate various Quality of Service (QoS) policies with each topic and dataflow between a DataWriter and DataReader. These include timing-related QoS that are implemented by the middleware. The Event thread manages both timing delays and periodic events such as protocol heartbeats, deadlines, and liveness needed to meet the QoS policies requested by the application. Database clean-up thread - This

thread purges old information from the internal data structures such as publication declarations and subscription requests. Receive threads – A port represents a transport-specific resource for receiving incoming messages. Data packets are delivered to transports ports. Different DataReaders can be configured to receive messages on different ports. In order to minimise the end-to-end latency, a receive thread is created per port provided by the transport.

When the application writes new data to a topic, the message passes all the way through the middleware down to transport level send, in the callers thread. In the users thread context, the message is serialized, deposited into the writer queue, encapsulated into a wire-protocol packet, and passed to the transport for delivery. In the common case, the entire operations critical path takes no inter-application locks and suffers no context switches. The event thread is only involved if the initial transport operation fails, or to execute follow-on processing such as ensuring reliable delivery. The event thread has ready access to the message, since it is already stored in the writer queue. When the transport receives a new packet, the appropriate receive thread processes the packet, retrieves the message, stores it in the reader queue, and immediately executes the listener callback. In the common-case critical path, there are no inter-application locks or context switches. If the application requires the message to be handled with user threads, it can do so with DDS WaitSets. Both flexibility and performance are optimised, even as the network scales.

Performance can also be impacted through the poor use of the code execution path. Since lock contention can have a significant detrimental impact on performance, fast path optimisation takes data to or from the network transport to the application using a single lock per message, greatly simplifying the resource-sharing protocol. Finally, instead of using lists to store the information needed to dispatch and manipulate messages, hash tables can be used. Although hash tables are more complex than lists, they have constant time access provided that the initial allocation of space is sufficient. Regardless, in the worst case, access time is logarithmic, which is better than linear linked lists.

This discussion shows just a few ways that an implementation of the DDS specification can impact the three critical characteristics of data communication over a distributed network – reliability, performance, and flexibility. Alternatively, a poor implementation of the DDS specification can mean that the architecture works well under certain optimal implementations, but fails to take advantage of greater resources, and fails to scale as the network grows.

Data communications system developers do not want to change their application code when the fabric is updated, changed, or augmented. However, many possible implementations can deliver suboptimal results when the network topology changes. A DDS implementation can take this into account so that the application can be easily re-optimised to deliver a comparable level of performance in the face of evolving and changing fabrics.

As switched fabric technology advances, the middleware must support those advances by being able to adapt to new transport mechanisms and different resource requirements and availability. Being able to plug in different transports in the middleware layer makes it possible to more easily incorporate new fabric technologies as they become available without making any changes at the application layer. A superior implementation of the DDS standard enables network performance to be optimised to the particular application. It matches the performance needs with the underlying fabrics and availability of system resources such as memory. The designs flexibility allows it to target a broad array of applications and network topologies by supporting many transports and maintaining individual resources for each connection. Finally, the design avoids most key single points of failure, increasing reliability. ■
