

## WHITEPAPER

# Is Your Data Secure?

## Executive Summary

Any organization's most valuable assets include its digital data. Safeguarding these assets raises questions: which assets require security, and what type of security is appropriate? Typical approaches for implicitly securing data include the standard user authentication, access control, integrity checks, non-repudiation, and confidentiality mechanisms at the domain and link layers. The challenge, and emerging requirement, is to attach security measures directly to the data itself. Security policies in this case must be targeted and managed at each element of system-identified data.

Why has data-specific security become a requirement? Consider the process of building out systems and systems of systems. How can security be managed when individual elements of data, the communication links, and domain boundaries have different related behaviors? Specifically, how can security be managed when connectivity does not necessitate access to all the data? With today's level of infrastructure and system complexity and the associated risks, making the right design and implementation decisions calls for increased and manageable security – at the data level.

In this white paper, learn about system deployment architectures, data models, security approaches, and the methods for more tightly aligning them. Learn how a Secure DDS solution can uniquely introduce fine-grained data security for data in motion. Review the features that make it possible to minimize the leak of information and any subsequent exploitation of data within a system architecture that is more scalable and flexible, and that performs optimally within stringent real-time and deterministic environments.

## Understanding Data and System Infrastructure

Too often the actual information and data that needs protecting is coupled with communication patterns, link protocols, and local and remote data-structures and storage. Fundamentally, it is the issue of what needs to be shared and distributed in a system versus how. Many existing security solutions and capabilities tackle the how. However, these solutions and capabilities also need to be orthogonally connected to securing the what.

An optimal security approach requires a full consideration of system requirements, performance needs, data criticality, link behaviors and properties, availability, and system risk. Since a security technique or process cannot address all issues, security must leverage best practices for any given system. A rigorous and consistent data security methodology requires decoupling the data from the protocol messages and handshakes, storage mechanisms, processes, applications and services.

### Data architecture

What is the data? Capturing knowledge of the system's information separately from its deployment and distribution requirements is critical for applying security mechanisms to the data. Essentially, a security solution can't protect undefined assets.

One approach that is gaining acceptance focuses on and documents the semantic interface between systems and components. This information must include the system's entities, the relationships between the entities

(context), and explicit definitions of the system's observable and measurable phenomena (such as attitude, distance, location, position, and so on). These definitions must include the reference frame, units, and precision to concretely establish the context of the data to be exchanged. The information, which could be in a model, must be rigorously defined, described, and discoverable.

Semantically aware software infrastructure, like RTI Connex, can then leverage this information and model for enhanced data awareness, security configurations, and implementations.

Once the data is known, a series of questions and system architecture concerns should be considered. Specifically, these items are concerned with form and behavior:

- What system topology is required to make the data accessible, actionable, and scalable?
- How is the information represented as digital values within the system?
- How can the data be best moved around? Who needs it, how much of it, and when?
  - Is the data at rest? On which servers? Hosted by which applications?
  - Is the data in motion? Over what connections and networks?
- What is the span of the data? Is the data shared?
  - Within a single group? Many groups?
  - Does everyone share a single view of state for the data or is state replicated?

These attributes of the data are not to be confused with the data itself. However, these attributes impact the approach used to make the data accessible, actionable, and securable.

## System architecture

Designers typically build a system around a deployment model that suits the system's data and the intended use cases. The resulting system architecture is then secured. This approach, however, couples the assumptions of the deployment and data access to a security model and is only appropriate if the scope, span, and use of that system and its data never change.

More typically, systems and systems of systems constantly evolve and change. In these cases, initial assumptions – about centralized storage, user authenticated access, performance requirements, presentation layers of data, and more – can quickly become invalid.

In general terms, there are two generic system architectures (Figure 1). Each satisfies different use cases and system requirements and most real systems blend the two. Whether centralized, distributed, or a hybrid, the deployment architecture often inadvertently dictates a security solution.




Figure 1: Centralized system architectures (left) offer simplicity for management of the information and data. Security concerns focus on the client, and its connection to the information, but centralized access can restrict performance and present challenges for scaling systems. Distributed systems (right), in comparison, make it easier to share and move data, but present challenges in authentication and access control. Security concerns here tend to focus on the perimeter and assume entities within the system are trusted.

Real systems (see Figures 2 and 3) exemplify different aspects of these two deployment patterns. Data that moves through these systems requires something more for security than just link or data-stores access-based security. There is no one protocol, central data-store, or even a common shared understanding of a user in the system. These heterogeneous systems require a heterogeneous security approach. The only common element is the data itself, which makes it logical to attach implementable security capabilities to the data itself.

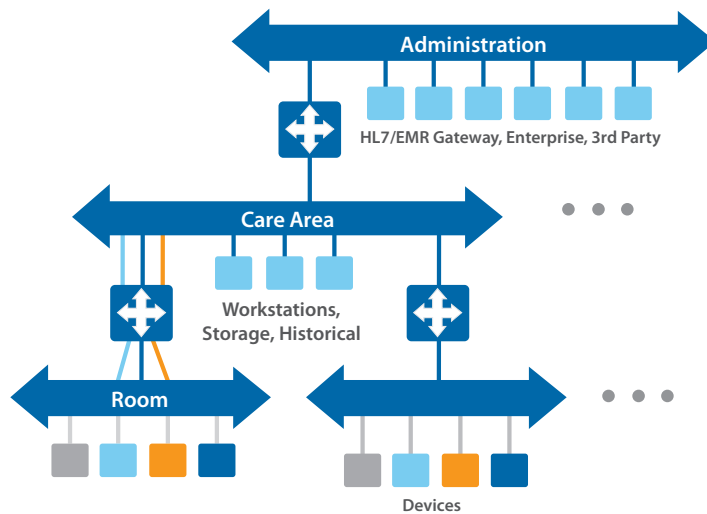


Figure 2. In a clinical decision support system, data is produced and collected from a variety of sensors and databases containing information about the patients. Patient sensor data is usually co-located with its displays and monitoring needs, but is also required to be connected to integrated healthcare services. For example, sensor data for a patient in room 100 is locally displayed, must be shared with the clinical team at the nursing station, and also be stored in the patient's electronic health records that are stored in the central data center. Data must also move with the patient as they move about the hospital. In this type of system, location and identify of the data is very important and needs to tightly coupled with the security implementations.

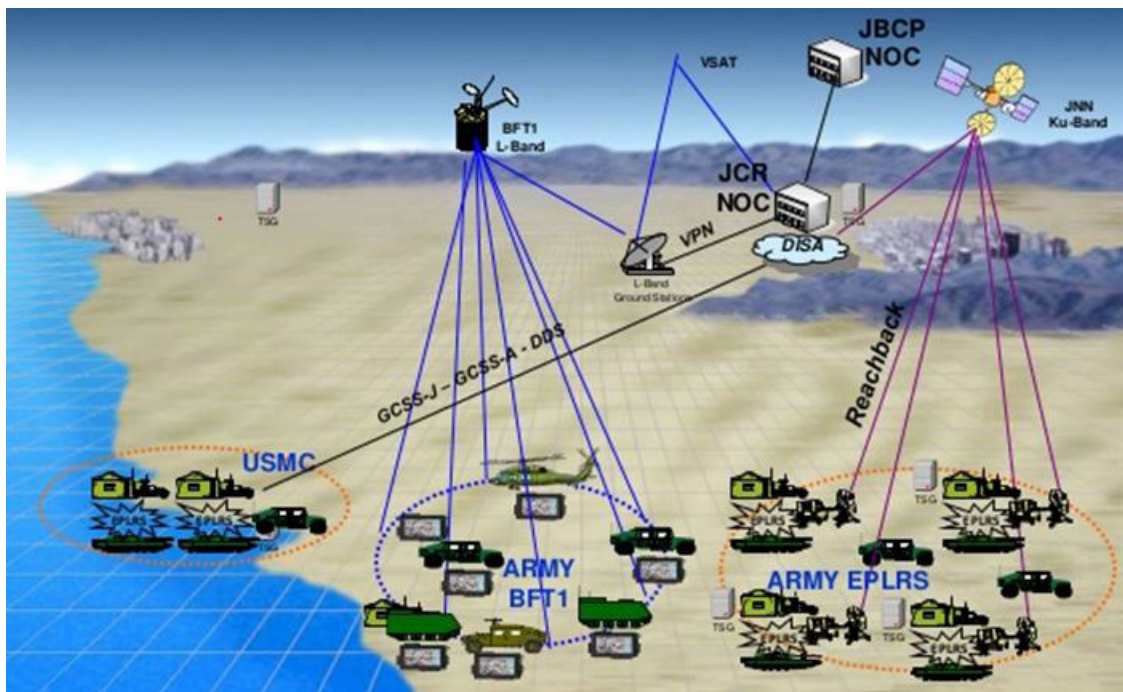


Figure 3. In a military tracking system, information is gathered in real-time to monitor and manage vehicle movements. On-vehicle devices collect and transmit position and other conditions. Observations from the troops are also collected within the system. The network operation centers (NOCs) aggregate the information, and maintain and disseminate a shared global picture. The connectivity complexity presents additional security challenges, and scalability of the overall system and security similarly pose complex requirements.

In these examples, the goal is to make the right data, at the right time, discoverable, actionable and available. This has enormous implications on security, which must be considered along with communication patterns, system topology, data format, and protocols.

Many systems today leverage point-to-point security protocols that provide authentication and access control at a link layer treating all data on the link equally. However, point-to-point connections that rely on message forwarding usually cannot scale for large, deterministic systems. Similarly, centralized approaches work well in the IT world and provide centralized security management capabilities, but limit the scalability and performance for complex deterministic and distributed systems.

In most systems, these implicit scalability constraints inherited from the security approaches are unnecessary and very limiting. Not all data is confidential. Some data just needs authenticated sources, some needs integrity checks, and other data just needs routeable confidentiality.

The many implementation and integration advantages make data aware, distributed systems the popular choice for large-scale systems. The challenge then becomes securing data that is now highly distributed and in motion. The overarching need is for a distributed, data-centric security architecture, which could overcome performance and reliability challenges, eliminate the need for a centralized data repository, and enable lower latencies and faster response times.

## Security Overview and Approaches


The security industry has evolved to encompass a broad range of methods and approaches for authentication mechanisms, access controls, confidentiality rules, and integrity checks and safeguards. Non-repudiation and availability solutions have also been introduced to ensure proof of access and actions within secure environments.

Best-of-breed security implementations similarly span a variety of security elements: intrusion detection and prevention, malware detection and blocking, secure boot and trusted platforms, secure communication links, key and identity management, cryptographic functions, and more.


These elements are all needed and form an integral part of a system's overall security profile. What's new is the addition of a data-centric scalable security capability.

### Data-Centric Security Boundary


Typically, security is system specific, with a different combination of security features applied at critical interfaces or boundaries. Now data can also be protected while in transit, at critical locations on networks, and protection can be attached to the data itself.



Boundary *security* locks down data by controlling access at a fixed entry point. Boundary security methods can include a locked door with armed guards, a computer that is not connected to a network, and a range of digital gateway security solutions that introduce safeguards at the main Internet or wide-area-network portal.



*Transport-level security* safeguards data in motion. These types of solutions span network security (layer 3 technology), session security (layer 2), as well as mechanisms on endpoints for encrypting and decrypting, for example. Virtual LANs, HTTPS, and other current infrastructure solutions employ transport-level security.



The most fine-grained approach, data-centric security, differs by separating security from the infrastructure and attaching it directly to each instance of the information. In the IT space, you see some examples for centralized data and secured access (role-based queue access, for example). In a fully distributed, peer-to-peer environment, the challenge involves natively securing data in motion, without the requirement for centralized management.

## A New Foundation for Aligning Security and Distributed Data

A new release of the Object Management Group (OMG) Data Distribution Service (DDS) interoperability standard defines an exciting new approach for securing data in motion. The latest specification builds on the DDS standard that was first published in 2004 which defines APIs for connecting applications to a data-centric global databus. In 2006, a wire protocol was added to the standard, facilitating vendor interoperability, highly available data paths, and consistent views of state under all conditions. Throughout the years, other enhancements addressed quality of service, extensible data types, fixed vs. dynamic access to data, and many other issues affecting interoperability.

The new DDS Security specification represents a major advancement for fine-grained data-centric security.

### DDS domains, participants, and topics

Logically, data (DDS Topics) within a system's network (DDS Domain) represents discreet instances of information that relate to the interest of the distributed applications and services (DDS Participants). Physically, the databus is instantiated by one or more DDS Participants creating publications (DDS Writers) and subscriptions (DDS Readers) of these discrete instances of data.

Participants include embedded devices and user-driven applications that are connected to the bus and functioning as sources of information or observers of information. For example, a heartbeat sensor in a hospital is a source of patient-critical data for one or more valid users.

Each sample (a discreet data value or an update of a previous value) on the bus is identified by a topic and unique identifier. Additionally, behaviors of the data are managed by DDS quality of service criteria (QoS) enabling highly configurable and data-specific delivery properties. Sensor data being streamed might be specified in terms of frequency and integrity while alarms can be specified in terms of reliability and availability. Quantity, history, durability (persistence) and liveness can also be managed and monitored via DDS QoS. One driving requirement for the new DDS Security specification was to remove limits on the use of these system and data-level QoS factors since this promotes the deployment of scalable systems.

### New DDS Security constructs

Starting with the inherent DDS capability for fine-grained data discovery and dissemination, combined with management of the acceptable data behaviors, DDS implementations can now be enhanced with new security mechanisms to offer data-centric security policies and controls. The DDS family of specifications has been expanded to offer the ability to secure a global data space with:

- Fine grained control of read and write access and permissions
- Peer-to-peer and serverless authentication capabilities
- Orthogonal data representation, transport, discovery, protocol, and security concerns
- Source-specific tagging
- Differentiation of read, write, and relay permissions

In general terms, these features allow control over who can access what information and who can control the dissemination of data for each topic. They directly map to the standard security concerns within the global data space:

- **Authentication:** User and device identities can be checked and validated as part of the process of granting access to a domain. Connection to the bus can be separately managed and specified from connection to the physical layers of the network.
- **Access control:** Even if connected, users must have specific read and write permissions for a specific topic or an identified set of topics.
- **Confidentiality:** Discreet data values and related metadata can be encrypted together or separately. Data payloads can be separated and encrypted for confidentiality while routing or in the case of implementing last-value historians.
- **Integrity:** Data samples can include destination-specific signatures or machine access controls (MACs). This can be applied independently or in parallel with confidentiality.

- **Non-repudiation:** Behaviors and associated quality of service can be specified for each source/sink of data to acknowledge receipt and acceptance of data to the level needed.
- **Availability:** Behavior management, fault monitoring, and failure actions can be designed into the domain on a per-data item basis.

## Building security into a DDS implementation

The new DDS Security specification spans four elements (see Figure 4). The Security Model documents the methods and approaches for mapping security policies and concerns to DDS actions, events, and APIs. The Security Plugin APIs specify how to instrument and connect a security library to core DDS functionality. This spans methodologies for peer-to-peer authentication and other security functions. Updates to the DDS RTPS wire protocol specify secure data encapsulation and discovery behaviors that maintain interoperability. And the default Built-in Plugin is specified to facilitate out-of-box implementations and multi-vendor interoperable implementations.

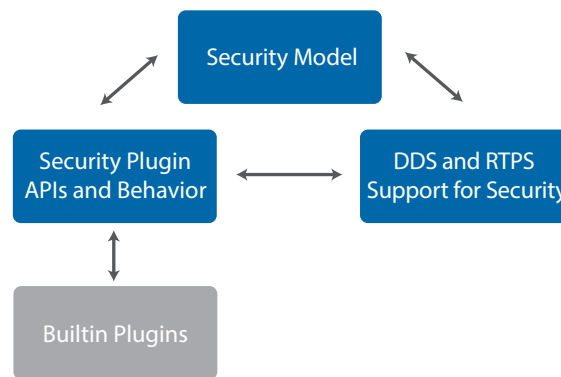


Figure 4. OMG DDS Security Specification

The specification does not require a central server or shared ubiquitous state. Applications can come and go. Regardless of the current system components, the specification fosters behaviors that are consistent and predictable within the global data space.

## Addressing threats within a data-centric, multicast-enabled system

The new DDS security mechanisms address the primary threats and define appropriate and inappropriate behaviors within a data-centric global data space. Appropriate behaviors on the databus include authorized subscribers reading only data relating to their approved topics, and authorized publishers only posting content on their approved topics.

Some participants (infrastructure services) might also be granted the ability to handle data without being able to read it. For example, an archiving service might be designed to store and replay data for certain topics without having the authorization to see the data content.

Additional concerns addressed by the specification include the ability to secure multicast environments. While most protocols such as Transport Layer Security (TLS) specifically preclude the use of multicast, the latest DDS specification allows for security concerns of access, confidentiality, and integrity to be independently configured and applied to each data element. This is critical when building systems with real-world performance and scale requirements.

For example, it might be required that sensor data be validated as coming from an authenticated source. Further, data integrity checks might be required to detect data tampering. However, even in cases that require these security measures, the data itself may not be confidential. If this data is also to be delivered with high availability, the underlying protocol of handshakes and heartbeats to ensure delivery and receipt must also be secure and robust to malicious attacks and manipulation. The new DDS Security solution addresses this type of complex situation and more.



The DDS Security specification is the first peer-to-peer, decoupled, multicast-capable solution. It is data aware, completely configurable, and enables the last piece in a secure system – managed security on the data itself.

## Configuring and Implementing Secure DDS

The Secure DDS specification facilitates this level of threat protection through the use of signed permissions documents that define access policies for each data domain. The configuration file lists the general access rules, discovery protection profiles, and the unsecured and secure topics for each domain.

For each secure topic within that domain, additional discovery authentication, confidentiality, and integrity mechanisms are defined to include encryption and/or signature-based controls. Further, the specification sets security policies for the metadata and payload data. This enables secure routing and storage of data samples, without exposing critical or private payload data to the system services.

Additionally, each DDS Participant (essentially an application on the databus) is configured with a signed permissions document that details the participant's identity and unique access privileges and rules for the identified data topics. Rules can be tied to individual participants, groups, or unique labels, and allow fine-grained control over allowable points of authentication and access. For example, the specification provides for control over which domains and topics each user (application) can access and in what roles (reader, writer, or relay), what partitions can be joined, and what labels and tags are associated with the readers and writers.

These separately signed permission documents dictate all security actions and policies within active DDS Domains. There is no requirement for centralized management of these documents and identity certificates and integration with PKCS11 capabilities, for example, is possible. The security plugin implementation is highly customizable to support best-of breed and system-specific security libraries, tools, and services. Essentially, the DDS security plugin instruments the standard dynamic DDS behaviors with checkpoints to enforce the configured levels of security.

The DDS protocol and wire activity specifications have been extended to support secure encapsulation of data. This encapsulation still allows all of the system scalability features of the Real-Time Publish Subscribe (RTPS) protocol, but adds fine-grained control of confidentiality and integrity.

## The Relevance of the New Security Paradigm

Secure DDS builds security into fully decentralized environments. As part of a DDS implementation, the security approach can be applied within a system regardless of the network transport (TCP, UDP, multicast, shared memory, etc.). In the case of real-time large-scale systems, multicast enables a highly scalable, low-latency architecture.

The enhanced DDS specification introduces very fine-grained control of access, confidentiality, and integrity. Encryption can therefore be employed only for data that absolutely requires maximum protection, without restricting the use of multicast and without imposing overhead by forcing encryption of all data. This results in secure systems that are significantly faster.

The data-centric security approach, based on customizable plugin architecture, delivers flexibility to the system architect. The fully DDS compliant foundation maximizes interoperability and works with unmodified existing applications.

## Incremental Adoption

The relevance and benefits of data-centric security can be illustrated within the example of a utility company's power distribution system. The company's existing grid included master and slave devices connected over point-to-point DNP3 links. Growing concern about the vulnerability of public power grids called for the introduction of attack and anomaly detection technology throughout the grid. However, the point-to-point topology made it impractical to introduce detection capabilities on every link and device.

With a Secure DDS solution, the company was able to re-instrument their cloud without changing devices, applications, or DNP3 connections. By leveraging existing data from all of the devices, and posting that information onto a global databus, the utility company made it possible to introduce centralized threat detection intelligence. See Figure 5.

Most importantly, this global integration was done securely. Each element had access only to what was authorized, and data was clearly identified on the integration databus. For example, by securing the data within the pipe instead of only the pipe itself, it was possible to offer appropriate and selective access for the Anomaly Detector.

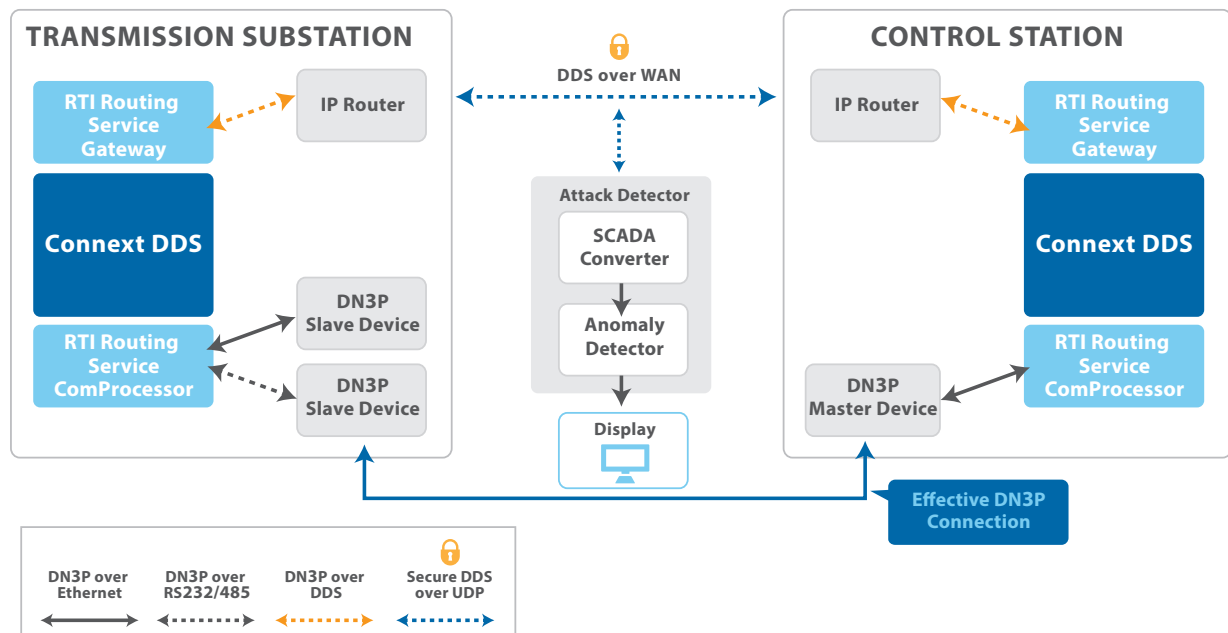


Figure 5. Retrofitting a power grid for security, and introducing threat detection solutions.

## Conclusion

Secure DDS enables secure, flexible, scalable, and high-performance integrated systems. The introduction of a data-centric security model offers many advantages compared to traditional boundary-level or transport-level approaches. It is possible to apply standard PKI and cryptographic techniques to enforce protection policies at the data level. Fine-grained controls can be attached to individual data elements, and configured uniquely for each domain, topic, participant, and location.

By focusing on the data rather than the boundary or transport layers, system architects can protect current investments and greatly simplify security-related application code.

## To Learn More

RTI Connex Secure DDS, the world's first turnkey DDS security platform, conforms to the OMG specification and provides a vital security infrastructure that is data-focused for DDS and legacy systems. Data-centric configuration policies make it possible to tailor security to a broad range of content and use cases, and a standards-based optional plugin SDK further enables the alignment of data security with system-specific tools and capabilities.

For more information about RTI Connex Secure DDS and how to design data security into your networks and connected systems, please visit the [www.rti.com](http://www.rti.com).

You can also download a free trial of RTI Connex DDS at <http://www.rti.com/downloads/>.



## About RTI

Real-Time Innovations (RTI) is the Industrial Internet of Things (IIoT) connectivity company. The RTI Connex<sup>®</sup> databus is a software framework that shares information in real time, making applications work together as one, integrated system. It connects across field, fog and cloud. Its reliability, security, performance and scalability are proven in the most demanding industrial systems. Deployed systems include medical devices and imaging; wind, hydro and solar power; autonomous planes, trains and cars; traffic control; Oil and Gas; robotics, ships and defense.

RTI is the largest vendor of products based on the Object Management Group (OMG) Data Distribution Service<sup>™</sup> (DDS) standard. RTI is privately held and headquartered in Sunnyvale, California.



Your systems. Working as one.

CORPORATE HEADQUARTERS  
232 E. Java Drive  
Sunnyvale, CA 94089

Tel: +1 (408) 990-7400  
Fax: +1 (408) 990-7402  
info@rti.com

[www.rti.com](http://www.rti.com)