Network Computing's storagepipeline Questionnaire

Dear Data/Information Lifecycle Management Solution Provider,

The following is a questionnaire for our use in understanding and evaluating the numerous solutions that are collectively called data or information lifecycle management solutions by vendors. Our working definition of these types of solutions is “a software and/or hardware solution for migrating data through the storage infrastructure using automated policies that match data content and/or access characteristics to storage platform capabilities and costs.” If you would like to suggest amendments or revisions to this definition, please make your suggestions as part of this questionnaire.

Please respond as per the cover letter attached to this questionnaire. You can edit this document directly to provide your information.

## Usage Scenarios

Since the capabilities of these solutions may vary, we are going to assess their effectiveness in the context of different usage scenarios.

**Capacity Utilization Efficiency.** These users mostly use a D/ILM solution to place data on platforms or media that are least costly, most capable and best suited to their access characteristics and “inherited” requirements in terms of retention, security, and criticality. Automated data migration may also be seen as a strategy for reducing TCO by reducing or capping administrative staff requirements.

**Capacity Allocation Efficiency.** These users are primarily interested in D/ILM as a mechanism for sharing capacity in the most efficient way and for eliminating unnecessary replication or junk data – often to defer additional hardware purchases. Automated capacity allocation may also be seen as a strategy for reducing TCO by reducing or capping administrative staff requirements.

**Disaster Recovery and Business Continuity.** These users seek to use a D/ILM solution to segregate data by its restoration priority in the wake of a disaster and to route data to the appropriate protection process (mirroring or continuous data protection, snapshot or disk-to-disk replication processes or tape backup). These users may also seek to leverage D/ILM to provide input to change management processes designed to keep plans up to date with burgeoning data.

## Description of Product

Troika has developed a complete set of hardware and software platform solutions, designed specifically to optimize the deployment of applications into storage networks. While Troika's products do not provide D/ILM capability directly, they are a necessary hardware building block in achieving desired ILM capabilities in next generation software. Troika's technology is being used by D/ILM software providers to host and provide data movement functionality for their products.

Troika's architecture is based on highly optimized, in-house ASIC technology purpose-built for optimizing low level software applications. This ASIC technology is augmented with a rich set of software components and middleware and an open application programming interface (API) carefully crafted to enable next generation storage software applications.

## Questions

**Development Status and Objectives.** Describe your vision of D/ILM and identify what features or functions your platform currently provides and what features or functions are still in development.

Troika provides a hardware and software platform for OEMs and ISVs implementing information lifecycle management solutions. Troika's application switch platform provides the ability to move data on a block-by-block granularity, as well as provide statistics on what data is being accessed most often by hosts. These functions are fully available in Troika's current generation of hardware, which has been generally available since May of 2003.

**Development Partners.** List the vendors with whom you are working to deliver your D/ILM solution functionality.

Troika is working with a number of Independent Software Vendors, including StoreAge, Incipient, and other undisclosed ISVs. Troika is also working with a number of tier-one OEMs but cannot discuss those relationships at this time.

**Data Naming Scheme.** What is your method for identifying data storage requirements or characteristics and for using those requirements or characteristics in building policies?

As an enabler for D/ILM, Troika does not provide any information for determining data storage requirements or characteristics. Our third party ISVs/OEMs do the discovery of these characteristics.

**Access Frequency.** Identify the mechanism by which your solution identifies how frequently specific files or datasets are being accessed and whether and how access frequency is used in migration policies.

Troika's application switches track statistics on which data is being accessed most often. These statistics are tracked by the Troika Application Switch hardware and firmware, without impacting application performance or data access speeds. This information is provided through programming interfaces to third party software, which can use this data to determine their own migration policies.

**Storage Platform Characterization.** Explain the mechanism that your solution provides for characterizing the performance capabilities and costs of specific hardware platforms for data storage so that this information can be used to target the appropriate storage platforms as destinations for automatically-migrated data.

Troika does not do any specific characterization of performance capabilities and cost; we rely on third party solutions to manage and discover this information. Troika's products may be used to do the actual migration and targeting, but is not aware of the specific characteristics of the storage.

**End of Useful Life.** Explain how your platform facilitates the automated removal and clean-up of data that has outlived its useful life and restoration of freed capacity for use by applications.

One of the base functions of Troika's application switches is providing block level storage virtualization. Block level storage virtualization helps with the D/ILM process by making it easier to allocate and free capacity for use by applications.

**Policy Articulation.** Describe how policies are created and how they are applied to existing data.

Troika does not provide any policy articulation, which is implemented by our partners and third party software vendors using our platform.

**Device Support.** What storage devices does your product support? Are their any proprietary devices (controllers, arrays, HBAs, switches, SAN topology, virtualization products, etc.) that are required for your solution to work?

Troika's products support a wide variety of Fibre Channel storage devices, HBAs, and switches. Because Troika's application switches are N-port based products (and look like a storage device or HBA) there is no issue with switch interoperability. Troika supports all of the major switch, HBA, and storage vendors, and most recently announced interoperability certification of Emulex Host Bus Adapters. Because of our fabric-neutral device, we also are able to connect to multiple, incompatible, heterogeneous storage area networks at the same time. One of the uses of our product is presenting the same storage to multiple fabric vendors at the same time (for example, presenting the same storage from one SAN to a Brocade, McData, and Cisco SAN at the same time).

**Resource Consumption.** How much bandwidth and server CPU “overhead” is introduced by your solution (e.g., to support polling processes, migration processes, agent processing, access frequency counting, etc.).

One of the key benefits of Troika's application switches is freeing server CPU cycles. Troika's high-powered application switches do all of the polling and processing and access frequency counting independently of the host CPU. In addition, hardware-assisted migration provides for very high performance, CPU-independent migration of data.

**Flexibility.** Does your solution have multi-vendor support? Can data, once integrated into your D/ILM scheme be migrated readily between your solution and other solutions in this space? Is you solution interoperable with other solutions in this space?

Troika is active on the T11 Fabric Applications Interface Standard (FAIS), which is focused on allowing vendors to easily develop and port their applications to

the various solutions in this space. Troika is an officer of the group, and has submitted proposed requirements for the FAIS specification to the group for approval, and is working actively with members to help make this standard a reality. Once implemented, the FAIS standard should allow D/ILM solution providers to more easily migrate their solutions to different platforms, including our own.

**Speeds and Feeds.** How should a prospective customer compare the relative performance of competitive solutions in this space? What are the appropriate performance measures or metrics to use in evaluating competitive solutions?

Troika measures its performance in the number of IOPS per port pair, and total throughput depending on block size. Our performance for a single port pair on our platforms is over 120,000 IOPs at full wire speed, and an aggregate box throughput of 460,000 with all 16 ports fully loaded. Troika's solution operates at full wire-speed, and operates at latency of an FC-switch.

**Interconnect Support.** List the network or fabric interconnects that you support. Is there an optimal interconnect for data movement for your solution? What about for management?

Troika's current solution provides for 8- or 16-ports of Fibre Channel connectivity. Troika's solution is fabric-neutral, and can also be used to span multiple, incompatible fabric networks. Troika's application switches do not participate as part of the Fibre Channel fabric but appear as hosts and storage to the network. Troika's application switches also include a 10/100 Ethernet port for web-based management. Troika has tested its products with all of the major fabric vendors in the space, including Brocade, McData, Qlogic, and Cisco. Troika's product can be managed either in-band through Fibre Channel (depending on the ISV software, which is hosted on our platform) or through Ethernet (for box and chassis management).

**Protocol Support.** List the storage networking technologies supported, including FCP, Ethernet, iSCSI, FCIP, iFCP, Parallel SCSI, SAS, and network file system protocols (NFS, CIFS, HTTP, DAFS).

Troika's products support FCP today and are expected to support iSCSI in the near future. Ethernet is supported as an external management interface.

**Data Type Support.** List the data types supported by your solution, including file types and file systems, databases, and hybrids such as email. Comment on the granularity of your D/ILM solution: will it support the migration of database components or subsets, subsets of email files, etc.?

Data type support is determined by third party provider. Some of our software providers are able to migrate very fine-grained data, on a block-by-block basis, enabling the migration of specific database tables or rows, and subsets of files at block granularity. Troika's hardware is capable of block level granularity, however implementation is dependent on the ISV's application.

**Cost.** List the cost of your solution or provide some means for calculating cost for a specific environment.

The list price of Troika's solutions are $28,245 for a 16-port solution, and $17,425 for an 8-port solution. Costs do not include optics or third party software that may be running on our platform.

**Standards.** List any relevant open standards upon which your product is built.

FC-PH
FC-AL
FCP-SCSI
SNMP
FibreAlliance MIB
SOAP/XML (external API access)
T11-FAIS

**Other features.** List other pertinent aspects of your solution.

**Miscellany:** Please note any additional information that you think would be worthwhile for prospective customers to consider about your solution or other solutions in this space.

## Brief Background on Troika's Technology

The following background might be useful in understanding our solution and how we fit into the I/DLM market as an enabling technology.

### The Need for Hardware Acceleration

As data movement applications have evolved in storage area networks, there has been an increase in the types and demand for more and more functionality from users. This increase in functionality has increased the amount and types of data movement required for software solutions.

Software providers are finding that running software alone on hosts or as appliances in the network is not sufficient for the robustness and performance these next generation applications require. As simpler applications such as virtualization, snapshots, and mirroring (all done both in hosts and appliances) move towards replication, migration, asynchronous replication, and beyond toward the holy grail of Information Lifecycle Management--the demand for data movement platforms has increased.

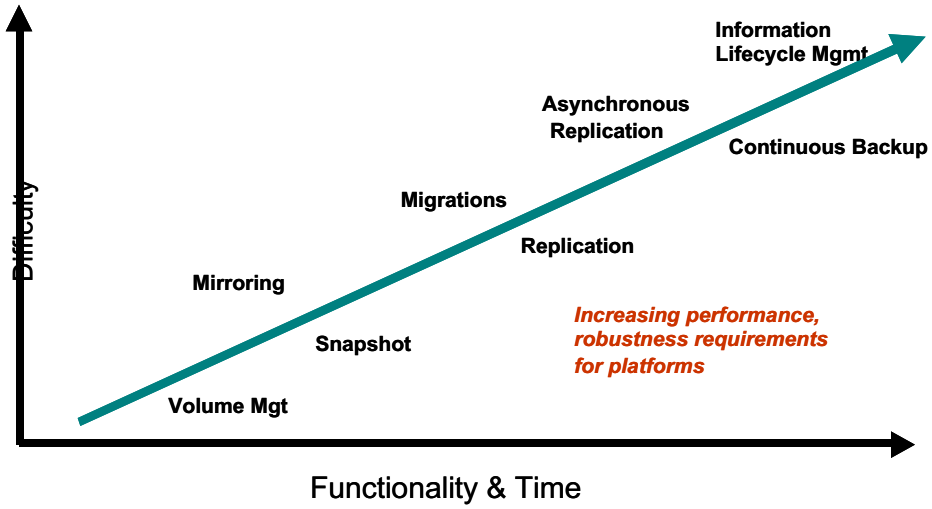


Software providers are finding that they can't achieve the performance levels their customers are demanding without taking advantage of accelerated data movement platforms.

## Accelerating Applications in Hardware

In the time honored tradition of accelerating software applications by moving software functionality into hardware, Troika has carefully moved key storage software functions into a highly integrated, cost effective ASIC platform. Troika's ASICs provide low level storage virtualization and related primitives in hardware, enabling software providers to deliver high performance, reliable solutions. Troika's software stack enables OEMs and ISVs to quickly and easily port their solutions to Troika's platforms and to take advantage of this hardware offload.

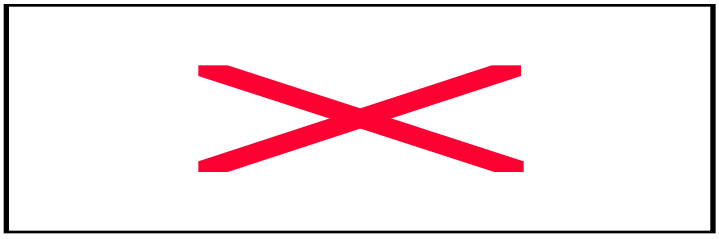

Figure 1. Split I/O, Control Handling

## Troika's Software Infrastructure

Troika's Software is not just an interface layer but a complete software infrastructure for storage applications. Troika's software stack provides a high level of functionality and abstraction of I/O primitives, allowing software providers to easily and quickly port and develop storage software applications to Troika's platforms.

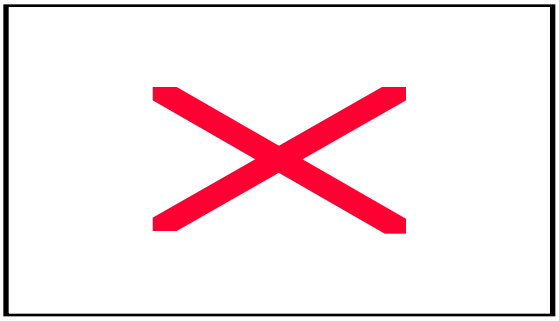

Figure 2. Troika's software stack provides a high level of functionality for software applications.

**N-Port, Network Storage Service Platform Approach**

Troika's application switches are implemented as Network Storage Services Platforms, which do not provide Fibre Channel switching but instead are Fibre Channel N-ports. This means that Troika application switches look like a storage array or HBA to the network, and do not participate in the switching. Troika's application switches provide I/O translation at switch speeds (a subtle but important difference).

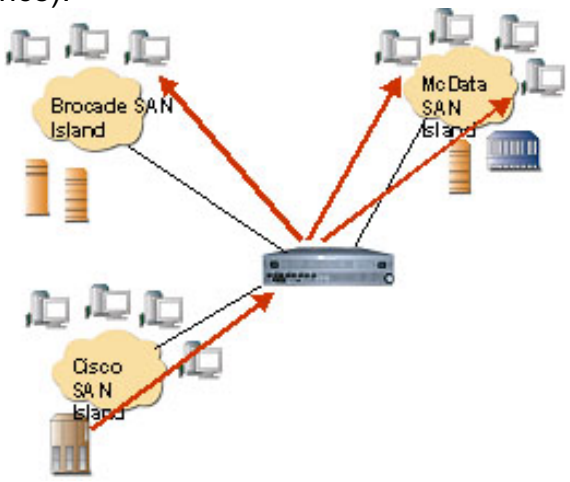


Figure 3. Example of exporting storage across incompatible fabrics

One of the advantages of this approach is the support for multiple, incompatible, heterogeneous Fibre Channel fabrics. As anyone in the SAN industry knows, despite the hype around interoperability of solutions there is still an issue connecting up disparate SAN islands.

This lack of interoperability has caused the proliferation of many, independent SAN islands, and prevented the full benefits of storage consolidation. Troika's approach allows for the bridging of these SANs, by allowing software providers to export storage through Troika's application switches (virtualized or non-virtualized) without connecting those fabrics. Each port of Troika's application switches can be independently connected to different kinds of fabrics. This isolates the issues with fabric compatibility and isolated fabric change events, and all of the issues associated with merging fabrics from different vendors. This also allows for a central place to run management software, where that management software can have visibility into multiple SAN islands.