



# Accelerating Data Protection

Enhancing Replication and Backups across Distance

---

## Executive Summary

This white paper looks at the challenges facing organizations as they struggle to ensure that critical data and systems are protected in the event of widespread disruption. Network replication and periodic backups between remote sites are two of the most common ways to provide Business Continuity and Disaster Recovery (BCDR). If one site goes down, the organization can continue doing business using the replicated systems and data at the other site. The scalability of this model is being tested, however, by the rate at which data is growing—burgeoning digital assets threaten to outpace the ability of replication and backup implementations to keep up. At the same time, the pressure for aggressive Recovery Point Objective (RPOs) and Recovery Time Objectives (RTOs) is not letting up.

Among the solutions available, WAN optimization seems to hold a lot of promise—at least on paper. Unfortunately, because traditional WAN optimization was designed for end-user applications and branch networks, it has not been able to scale up to the levels necessary to ensure current and future data protection. To address the need for high-performance WAN acceleration, Infineta Systems has developed a set of new technologies that are specifically designed to improve the performance of inter-data center applications, including replication and backup. The Infineta solutions work at speeds of up to 10 Gbps and combine data reduction, TCP optimization, and Layer 4 QoS to allow organizations to move more data across the same WAN infrastructure in less time.

---

## The Provenance of a Good Business Practice

More than 4,500 years ago, the Phoenicians divided their inventory among multiple vessels when transporting goods along the Mediterranean Sea, lest one be lost in a storm or to pirates. Today, most organizations follow the same practice, although now, in addition to physical goods, inventories include millions or billions of dollars in digital assets. These assets are no less vulnerable to the disruptions and loss that threatened the Phoenicians—acts of nature, pirates (hackers), social upheavals, and war. In fact, the risks facing digital assets are greater, because, the data is just a complex arrangement of magnetic polarities on a thin electromagnetic film. Anything can happen, at any time.

The good news is that protecting digital assets only involves making a copy and storing it in a different location than the original. This is fast and easy—at least for one document or one user. On the scale of an entire organization, however, where data is being measured in petabytes, soon to be exabytes<sup>1</sup>, protecting a multi-billion dollar digital inventory is a daunting task. The responsibility for BCDR may ultimately rest with the Board of Directors or CEO, but the onus usually falls on the Chief Information Officer, Chief Information Security Officer, or Chief Security Officer. While the goal is typically to ensure continued access to critical systems and provide at least some level of recoverability, the question of whether operations can be restored in minutes, hours, days or weeks depends on the quality of the BCDR plan and how well it can be carried out.

---

<sup>1</sup> In November 17, 2010, Chuck Hollis, CTO EMC Corporation, noted in his blog that EMC has 1000+ customers in its “Petabyte Club.” On January 14, 2011, he forecast an “Exabyte Club,” likely in 2012, for EMC customers who have reached 1000 petabytes of storage. [http://chucksblog.emc.com/chucks\\_blog/2010/11/big-data-and-the-ever-expanding-emc-petabyte-club.html](http://chucksblog.emc.com/chucks_blog/2010/11/big-data-and-the-ever-expanding-emc-petabyte-club.html)

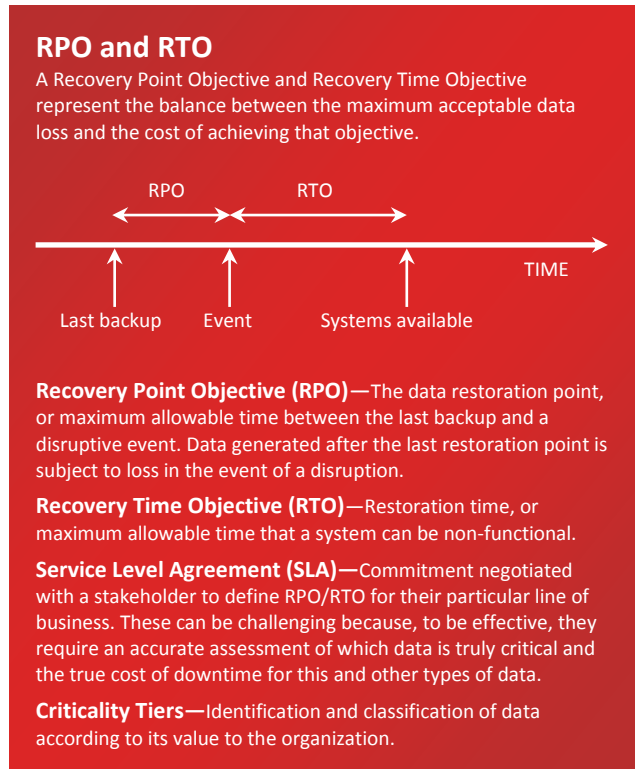
## Classifying Data by Value

To protect digital assets, the IT organization, in partnership with the relevant lines of business (LOB), determine the value, usually according to the following priorities:

- How much of the asset the organization is willing to lose
- How long the organization can go without the asset

The answer of each of these is represented by metrics known as the Recovery Point Objective (RPO) and Recovery Time Objective (RTO). The IT organization is usually responsible for meeting the RPO and RTO, while at the same time managing costs, which is then expressed through service-level agreements (SLAs) with the LOBs. In general, the lower the RPO/RTO requirements, the greater the cost of meeting them. Setting the RPO and RTO to zero for all assets and all LOBs would be ideal—no data lost, no service interruptions, and no complaints—however the expense and logistics of doing so can be unwieldy.

To balance the cost of protection against the value of assets, the assets may be classified into criticality tiers. For example, a national retailer that records millions of transactions each month may determine that revenue-generating systems must always be available, that is, no transactions can be lost even in the event of a major disruption. Here, sales data and systems would be classified as Tier 1, with a service level agreement promising an RPO and RTO of zero. The implication is that all the transactional data and systems will simultaneously exist in two or more places. If any system, storage device, or even a whole site is compromised, the business will fail-over seamlessly to the other system, data, or location. High costs typically mean that only the most critical data is classified as Tier 1. Less critical assets are classified as Tier 2 or Tier 3, but this can be difficult negotiation.



## Growing Storage Requirements

Organizations today are committed to digitalized data, and this is having a profound effect on storage and transport. Historically, businesses decided for themselves which assets they needed to protect and the choice was based on factors such as convenience, cost, and tolerance for risk. Their only external obligations were to meet the expectations of customers, partners, and shareholders. More recently, however, legislation has played an increasing role in regulating storage practices by defining which data an organization must retain, for how long, and under what conditions.<sup>2</sup> For the financial industry, the

<sup>2</sup> Additional statutes and stipulations in the USA regarding data protection and retention include, FINRA (Financial Industry Regulatory Authority, securities regulations), 21 CFR Part 11 (the US Food and Drug Administration, record retention), the Gramm-Leach-Bliley Act (requirements for securing nonpublic consumer information), FRCP (Federal Rules of Civil Procedure, e-

impetus has been post-September 11 economic concerns regarding the interdependence of major banking institutions.<sup>3</sup>

Among storage analysts and professionals, it is a commonly held belief that for most large organizations, the rate of data storage is doubling every three to four years, and the amount of WAN traffic being transmitted between data centers is doubling every two years.<sup>4</sup> Because inter-data center traffic is comprised primarily of BCDR traffic

(i.e., Tier 1 and Tier 2 assets), expanding storage demands are challenging the ability of organizations to meet existing RPO/RTO commitments. The crux of the matter is that as the volume of critical data grows, with RPO/RTO goals staying the same, the “pipe” between data centers has to get bigger, fuller, or faster to handle the increases. The other alternative, of course, is to leave the pipe as is and protect less data. The

most common solution thus far has been to enlarge the pipe through a series of bandwidth upgrades. As Figure 1 shows, however, this method cannot be expected continue as storage (and the volume of Tier 1 data therein) grows from petabytes to exabytes in the coming years.

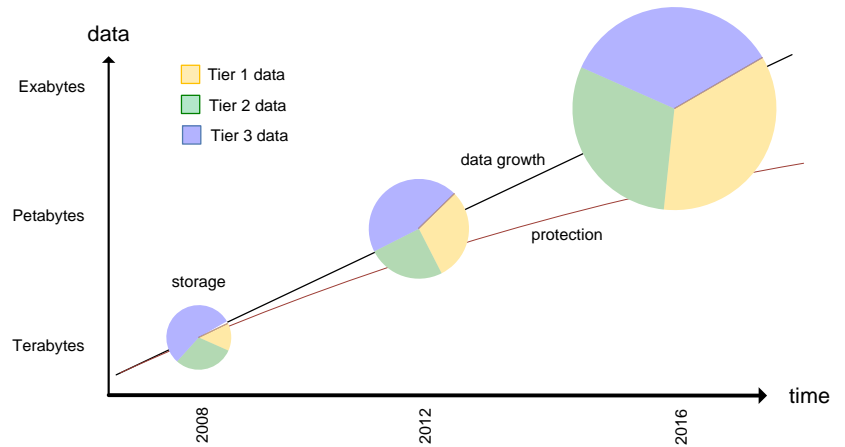


Figure 1. As storage demands grow, so does the amount of data requiring protection, which makes it increasingly difficult to meet the same RTO/RPO goals.

**Next...** The following sections looks at the most widely used methods of copying critical assets from one location to another, and show why bandwidth is inherently limited as a means of keeping pace with growing storage demands. The paper then concludes with a look at how organizations can optimize their existing WAN links so more data can be transferred over the existing infrastructure, i.e., how the pipe can be made fuller, and faster.

## Protecting Digital Assets

When organizations have a lot of digitized assets to protect, the most widely used methods are replication and backup. The main differentiators between the two are time and expense, although they are often used in conjunction with one another to balance scope and cost. Replication is more costly to support, but it can occur in real-time or near real-time and provide lower RPOs and RTOs. Backups usually occur on a periodic basis. They have lower operations costs, but are associated with higher RPOs and RTOs.

---

Discovery), HIPAA (Health Insurance Portability and Accountability Act, data privacy), Sarbanes-Oxley (record retention and audit trails), PCI DDS (Payment Card Industry Data Security Standard, retention and privacy).

<sup>3</sup> In 2003, the Federal Reserve Board, the Office of the Comptroller of the Currency (OCC) and the Securities and Exchange Commission (SEC), published guidelines that focus on minimizing the immediate systemic effects of a wide-scale disruption on critical financial markets (“Sound Practices to Strengthen the Resilience of the U.S. Financial System”). In 2010, the United States Department of Homeland Security finalized its Voluntary Private Sector Preparedness Accreditation and Certification standards, and the governments of Australia and Great Britain produced similar documents.

<sup>4</sup> Forrester Research, Inc. May 2010. “The Future of Data Center Wide-Area Networking.”

## Replication

Replication can be synchronous, which means each byte of data is written to multiple locations before the application can write new data, or it can be asynchronous, which means that remote store will always be some amount of time (and thus data) behind the source. Because of its sensitivity to time, synchronous replication is usually restricted to sites within 100 KM of each other,<sup>5</sup> whereas asynchronous replication can occur at any distance and so is more common. As the distance between asynchronous sites grows, the data discrepancy between them also grows as a result of the increased latency. The difference can be anywhere from a few writes to many gigabytes of data, and at some point the difference between protecting data through replication or backups becomes indistinguishable.

Technique	SLA	Characteristics
<b>Replication</b>		
Synchronous	RPO = 0 RTO = 0	<ul style="list-style-type: none"> <li>• Bandwidth intensive</li> <li>• Distance/latency limits</li> </ul>
Asynchronous	RPO = minutes to hours RTO = minutes to days	<ul style="list-style-type: none"> <li>• Bandwidth intensive</li> <li>• Calculated risk of data loss</li> </ul>
<b>Backup</b>		
Continuous	RPO = minutes to hours RTO = hours to days	<ul style="list-style-type: none"> <li>• Storage intensive</li> <li>• More data in jeopardy</li> <li>• Higher RTO</li> </ul>
Snapshot	RPO = hours to days RTO = hours to days	<ul style="list-style-type: none"> <li>• Storage intensive</li> <li>• Designed around restore points</li> <li>• Higher RTO</li> </ul>

Table 1. Replication and Backups are used for data protection.

Replication flows between data centers are not like branch-to-data-center traffic, which is usually comprised of many small, short-lived, low RTT connections. Replication flows tend to have unique characteristics and require specialized resources to overcome the limitations imposed on it by the WAN.

- **High speed**—Connection speeds can be as high as 1 Gbps per connection.
- **High volume**—Replication traffic is constant and can total terabytes each day.
- **Few connections**—Replication traffic uses relatively few connections compared to typical end-user scenarios (tens or hundreds of connections vs thousands or tens of thousands).
- **Long-lived connections**—Replication connections are persistent, lasting days or months, while application delivery connections are often created on a per-transaction basis.
- **Latency sensitive**—Replication traffic is highly sensitive to latency.
- **“Bursty”**—Data transmissions start and stop suddenly and frequently.

<sup>5</sup> Synchronous replication requires that new data be written and confirmed in two locations before the next new data will be accepted. The total latency budget—from endpoint to endpoint—for synchronous replication usually cannot exceed five milliseconds, i.e., a distance of roughly 100 KM.

## Multi-hop replication

Some organizations have found another path forward by using a mix of synchronous and asynchronous replication in conjunction with multiple levels of backup, the so-called multi-hop strategy. Figure 3

illustrates how multi-hop replication can be used to work-around to the limits imposed by latency.

Synchronous replication is used to ensure that Tier 1 data exists in multiple locations at the same time.

Asynchronous replication is then used to copy the data from the intermediary site to a more distant location.

Backup applications can also be employed to capture and move snapshots of the remaining data to remote locations to ensure that Business Continence Volumes (BCVs) are available for all data tiers. Although multi-hop seems like a workable, short-term solution to the problem of keeping up with data growth, it is equipment-intensive and scalability is not clear.

Solution-specific idiosyncrasies can also make managing the arrangement a challenge for IT staff.

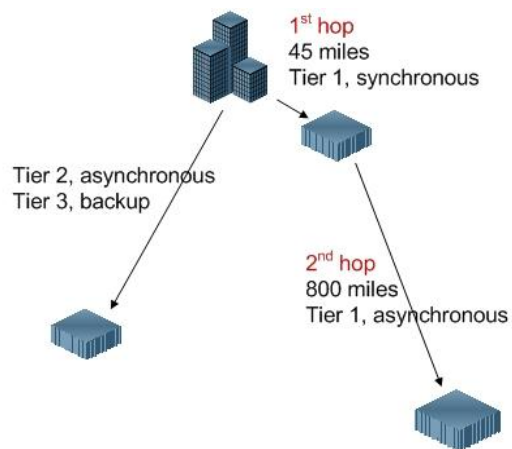


Figure 3. Multi-hop replication and blended protection have been used to reconcile the needs for geographic dispersion between data centers and zero-data-loss RTOs.

## Backups

Backup strategies are used to create a restore point for critical applications and data sets. In the event of disruption, any data generated after that point is subject to loss. For decades, backups were made to on-site tape drives, and the tapes were then hand-carried to an offsite location for storage. Tape backup tends to be slow, delivery to the off-site location is often “best-effort” basis, and the tape media itself may not support the lifespan of the data being stored on it. These issues, along with aggressive RPOs/RTOs, more Tier 1 systems and data, and improved Internet connections, have driven most organizations to turn to network-based backups, which include data snapshots, Continuous Data Protection (CDP), and periodic backups.

Vendor	Synchronous Replication	Asynchronous Replication	Backup
EMC	SRDF/S RecoverPoint	SRDF/A RecoverPoint	SRDF/DM
HDS	True Copy Synchronous	True Copy Extended Distance	Hitachi Dynamic Replicator
IBM	Metro Mirror (PPRC, or Peer-to-Peer Remote Copy)	Global Copy	Global Mirror
HP	HP Continuous Access	HP Continuous Access	HP Continuous Access HP StorageWorks Enterprise Backup Solutions (ESB)
NetApp	SnapMirror	SnapMirror	SnapProtect
Dell		CommVault Simpana Replication	CommVault Simpana SnapProtect
Symantec	Veritas Volume Replicator	Veritas Volume Replicator	NetBackup

Table 2. The Infineta DMS optimizes the WAN to accelerate high-speed data protection solutions.

Each of the network-based options has its own advantages and disadvantages, but the one thing they all have in common is that the scope of their protection is ultimately constrained by the WAN. As the amount of data being stored grows, so too do backup requirements, until at some point, the limit for how much data can be backed up within the available time frame (called the backup window) is reached.<sup>6</sup> Either the last backup will not be able to finish by the time the next one is set to begin, or the amount of data for backup will have to be reduced to fit the available backup window.

## The TCP Limitation

From a BCDR perspective, locating one data center in say, New York City and another in San Francisco would be ideal because whatever caused one site to go down is unlikely to also affect the other. One seldom sees collaborating data centers so far apart, however, because of the impact of latency on TCP throughput. Even at the speed of light, it takes time for data bits to move from one location to another, about 36 milliseconds in the case of NYC and SF. Yet TCP is widely used for critical data transports because it is connection-oriented, which means the receiving end of the connection periodically tells the sender that the packets it has been sending have been received. Confirmation adds another 36ms to the packet delivery time, for a total round trip time (RTT) of 72ms. So although TCP provides the benefits of assured delivery and error recovery (for packet loss), they come at the cost of throughput, i.e., the rate at which the data can be *successfully* sent from one site to the other.

Figure 2 shows the relationship between RTT and throughput. As noted above, the greater the distance between sites, the greater the RTT. With TCP, delivery of new packets is held up until confirmation of old packets is received, which limits connection throughput. Thus, even on the 1Gbps WAN link, the

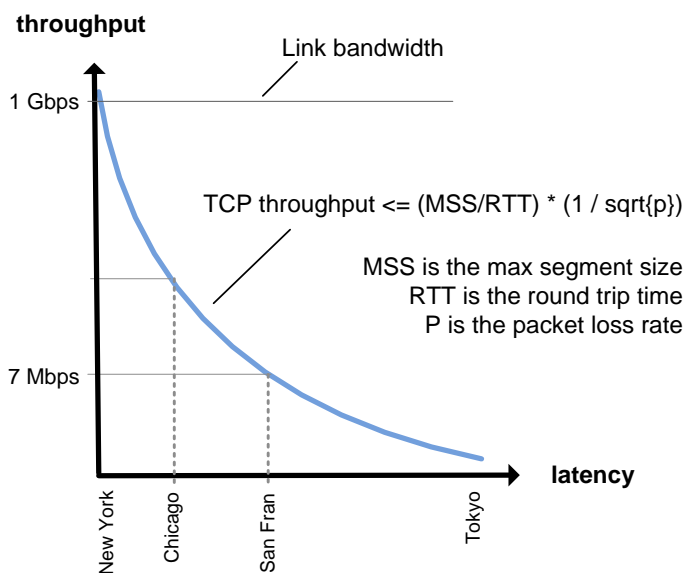


Figure 2. Replication and Backup traffic rarely achieve fully bandwidth utilization on long-distance links. Although bandwidth upgrades may improve throughput, the efficacy is subject to diminishing returns.

maximum throughput for a single TCP connection between New York and San Francisco is about 7Mbps—not 1Gbps as might be expected. Adding bandwidth to the link cannot increase the per-connection throughput. The impact of latency on TCP throughput is particularly acute with regard to replication flows, because as we have seen, these typically involve a relatively small number of long-lived connections. Increasing bandwidth may provide marginal gains by providing capacity for new connections, but to date, the only real recourse for the issue of latency is to reduce it, usually by locating data centers in closer proximity to one another than might be considered ideal when planning for BCDR.

<sup>6</sup>To keep up with daily changes of 5% to 2 petabytes of storage would require more than a day: Two petabytes = 2,000,000,000,000 bytes. Five percent of 2 trillion = 100GB, or 800 gigabits of data. Transferring 800Gb at 7Gbps, the throughput between NYC and SF, would take about 32 hours.

## Traditional WAN optimization

WAN optimization has emerged over the past decade to reduce the effect of latency on end-user applications between the branch and data center. But the original technologies were designed to improve the performance of chatty branch applications and the access of oft-requested objects (e.g. files and email attachments). So, while effective at reducing end-user wait times when accessing applications remotely, they are unable to scale to the multi-gigabit flows necessary for inter-data center traffic. Ultimately, the deduplication algorithms at the core of the technology are too CPU intensive to process traffic at the rates required for inter-data center WAN flows—the ratio between CPU cycles and unit of reduction is too high to be used with latency-sensitive traffic or on long-distance, high-speed WAN links. These “legacy” solutions are best suited for branch WAN scenarios, i.e., where RTTs are low, connections are small, and the impact of imposed latency is not critical.

## Accelerating Data Protection

Businesses today require a cost-effective way to protect their current investment in digital assets. That solution must also be able to provide a path forward so they can continue to meet RTO/RPO, even as the storage repositories grow. Such a solution should not, however, require yet another round of upgrades to the WAN infrastructure just to keep from falling behind.

Infineta Systems has developed a seminal new approach to network deduplication that makes data reduction possible for WAN traffic at speeds of up to 10 Gbps. It is called the Infineta Data Mobility Switch (DMS) and enables organizations to replicate and/or backup as much as 20 times the data previously possible, without needing to upgrade WAN bandwidth or data center infrastructure.


### Designed for Speed

The DMS is designed especially for latency sensitive, high-speed traffic and is compatible with all leading data protection solutions. It works by removing redundant byte patterns from replication, backup, and other inter-data center network flows, which reduces the footprint of WAN traffic by 80% to 90%, even at speeds 10 Gbps. At the same time, the DMS addresses the long-standing problem of TCP performance on high RTT links. It then utilizes the new capacity gained from the reduction by increasing the connection speed of existing flows, which in turn, are reduced, and so on until the link is fully utilized. In this way, end-to-end throughput can achieve the full bandwidth speed of the link. Backups complete in a fraction of the time previously required, and vastly more Tier 1 and Tier 2 data can be replicated in the same time.

In the New York to San Francisco example (RTT of 72 ms), where the maximum per-connection throughput on a WAN link between sites was 7Mbps, adding a DMS to each end of the WAN link would immediately increase throughput to about 1 Gbps for a single connection. Combined with reduction provided by the same DMS, an optimized 1Gbps link could carry the equivalent of 10Gbps data. In other words, applications that were constrained to 7Mbps per connection due to the WAN could start performing at speeds of 10Gbps given TCP optimization, which fills the 1Gbps link by mitigating the effects of latency, and reduction, which reduces the size of the data footprint by as much as 85%.

The DMS is specifically designed to meet the unique demands of high-performance, low latency traffic such as occurs with replication and backup workflows. It creates a fuller, faster pipe for data transfers and provides an alternative to the cycle of upgrading WAN bandwidth to keep pace with growing storage. In short, the DMS provides:

- **High per-connection throughput**—Accelerates a single TCP connection all the way up to 1 Gbps, which is critical for the traffic bursts of modern replication workflows.
- **Optimization for high-latency WAN links**—Fills the existing WAN link regardless of distance by employing highly aggressive start and recovery algorithms
- **Very low port-to-port latency**—Averages 50 *microseconds* latency, which means acceleration for both synchronous and asynchronous traffic between data centers
- **Complete protection against packet loss**—Protects TCP connection speeds by handling packet loss more efficiently than could the endpoint
- **Packet ordering**—Ensures correct packet order delivery to the endpoint to eliminates the need for retransmits



**Figure 4.** Optimizing Tier 1 data lowers RPO/RTO and bandwidth requirements. Optimizing Tier 2 data provides flexibility for BCDR planning by allowing more distance between sites and reducing WAN costs.

## Agile, Future-Proof WAN Infrastructure

Once the constraints of distance and speed have been alleviated for replication and backup in the near-term, organizations can begin working on strategic IT initiatives such as developing active/active data centers, implementing cross-site virtualization, or processing Big Data across multi-site clusters.

Whatever the goals of your organization turn out to be, Infineta's DMS can play an enabling role in the transformation of long-distance WAN links from a limitation that must be overcome, into a platform that is agile and prepared for the future.

## About Infineta Systems

Based in San Jose, California, Infineta Systems is a privately held provider of WAN optimization systems for Big Traffic. The company's patent-pending Velocity Dedupe Engine™ delivers unprecedented levels of throughput, scalability and bandwidth capacity to support critical machine-scale workflows across the data center interconnect. Its flagship product, the Infineta Data Mobility Switch, accelerates multi-gigabit BCDR (Business Continuity/Disaster Recovery), cross-site virtualization, and Big Traffic. The company is backed by Rembrandt Venture Partners, Alloy Ventures and North Bridge Venture Partners.

For more information, visit [www.infineta.com](http://www.infineta.com).

### Contact Information

2870 Zanker Road, Suite 200  
San Jose, CA 95134

Phone: (408) 514-6600

Sales / Customer Support: (866) 635-4049

Fax: (408) 514-6650

General inquiries: [info@infineta.com](mailto:info@infineta.com)

Sales inquiries: [sales@infineta.com](mailto:sales@infineta.com)

©2012 Infineta Systems, Inc. All rights reserved. No portion of the Documentation may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written permission of Infineta. Infineta disclaims all responsibility for any typographical, technical, or other inaccuracies, errors, or omissions in the Documentation. Infineta reserves the right, but has no obligation, to modify, update, or otherwise change the Documentation from time to time. *Infineta*, *Infineta Systems*, *Data Mobility Switch*, and *Velocity Dedupe Engine* are trademarks or registered trademarks of *Infineta Systems, Inc.*, in the U.S. All other product and company names herein may be trademarks of their respective owners.