

WHY SECURE MULTI-TENANCY WITH DATA DOMAIN SYSTEMS?

The Business Value of secure multi-tenancy for Data Domain

WHY YOU SHOULD TAKE THE TIME TO READ THIS PAPER

- **Provide data isolation by tenant** (Secure logical data isolation for each tenant by department or customer on a shared Data Domain system.)
- **Enable Data Protection as a Service with protection storage** (Improve operational efficiency for the provider while enhancing the tenant experience.)
- **Reduce service requests time** (Reduce bottlenecks to customer service requests through tenant self-service.)
- **Reduce backup storage costs** (More efficient utilization of Data Domain system resources.)
- **Enable tenant self-service** (Tenants are able to perform their own administration and reporting for their own environments.)
- **Manage Data Domain system resources per tenant** (Data Domain Admin controls storage capacity used and stream counts per tenant.)
- **Report on physical capacity used per tenant** (physical capacity measurement enables chargeback/billing, and information for capacity and migration planning.)

October, 2016

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA, 10/16, Business Value White Paper, H12937.4

Dell EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

TABLE OF CONTENTS

- EXECUTIVE SUMMARY4**
 - Secure multi-tenancy for Data Domain..... 4
 - Introduction..... 4
 - Audience..... 4
- DATA DOMAIN SYSTEMS DIFFERENTIATION5**
- INTRODUCTION TO SECURE MULTI-TENANCY5**
- WHY DATA DOMAIN FOR MULTIPLE TENANT ENVIRONMENTS?.....6**
 - Data isolation 6
 - Administrative isolation 6
 - Benefits of administrative isolation 7
 - Data path isolation 7
 - Benefits of data path isolation 7
 - Management..... 8
 - Tenant onboarding and provisioning 8
 - Benefits of tenant onboarding and provisioning..... 9
 - Tenant self-service..... 9
 - Benefits of tenant self-service..... 9
 - Metering and reporting 10
 - Logical capacity reporting 10
 - Physical capacity measurement..... 10
 - Alert monitoring per tenant 11
 - Historical streams performance per tenant 12
 - DD Boost stream usage per tenant 12
 - File replication reporting per tenant..... 12
 - Benefits of multi-tenant reporting metrics..... 13
- CONCLUSION14**

EXECUTIVE SUMMARY

SECURE MULTI-TENANCY FOR DATA DOMAIN

Multi-tenancy refers to the hosting of an IT infrastructure by an internal IT department or an external service provider for more than one consumer/workload (business unit/department/tenant) at the same time. A tenant is a consumer (business unit/department/customer) who maintains a persistent presence in a hosted environment.

Dell EMC's implementation of secure multi-tenancy for Data Domain systems provides the ability to securely isolate many users and workloads on a shared system such that the activities of one tenant are not visible or apparent to other tenants. This capability improves cost efficiencies through a shared infrastructure while providing each tenant with the same visibility, isolation, and control that they would have with their own stand-alone Data Domain systems.

INTRODUCTION

This paper focuses on Data Domain systems technology leadership and differentiation and why it matters to you. The purpose of this paper is to explore the technical and financial reasons why Data Domain systems are ideal for backup and archiving in a multiple tenant environment. With the introduction of secure multi-tenancy, Data Domain systems allow enterprises to deliver protection storage as a service in their private cloud environments and service providers in hybrid or public cloud environments.

AUDIENCE

This paper is intended for Dell EMC customers, Dell EMC sales, Dell EMC systems engineers, Dell EMC partners, and anyone else who is interested in learning more about Data Domain system's differentiating technology and all the unique advantages that it can provide for your backup and archive data.

DATA DOMAIN SYSTEMS DIFFERENTIATION

Dell EMC continues to lead the industry in Purpose Built Backup Appliances with 61.4% market share according to IDC as of Q2 2016. There are many reasons for this long-standing market dominance. This paper will focus specifically on the secure multi-tenancy feature for Data Domain. At the end of this paper you will find a number of links to additional documents that will explain in more detail why Data Domain systems have been so widely accepted. For the purpose of introduction and as a foundation to the rest of this paper, here are just a few of the reasons why customers choose Data Domain systems more than any other protection storage:

- **Data Domain Data Invulnerability Architecture** – End to End data integrity protection of your critical data in the storage of last resort better than any other solution
- **Deduplication efficiency** – variable length deduplication that typically provides between 10-30 x storage reduction for backup data
- **High-speed deduplication** – up to 68 TB/hour backup performance
- **Highly scalable** – protect up to 150 PB of data with a single system
- **Backup and Archive consolidation** – efficient storage of both and supports critical compliance regulations such as SEC 17a-4(f)
- **Network bandwidth efficiency** – use up to 99% less bandwidth
- **Fast, efficient and scalable disaster recovery** – thanks to efficient and secure deduplicated replication, Data Domain systems provide improved and more cost effective disaster recovery
- **Seamless integration** – with simultaneous support for NFS, CIFS, VTL, and DD Boost over Ethernet or FC, Data Domain systems work with all major backup and archive applications as well as integrating directly with enterprise applications such as Oracle RMAN, Microsoft SQL, IBM DB2, SAP HANA, and SAP with Oracle

INTRODUCTION TO SECURE MULTI-TENANCY

Secure multi-tenancy on Data Domain allows enterprises and service providers to deliver data protection as a service as illustrated in Figure 1. Specifically this feature:

- Enables enterprises to deploy Data Domain systems in a private cloud
- Enables service providers to deploy Data Domain systems in a hybrid/public cloud
- Allows for different cloud models for protection storage
 - Local Backup, Backup as a Service for hosted applications
 - Replicated Backup, Replication as a Service, DR as a Service
 - Remote Backup, Backup as a Service over WAN

Cloud Model	Large Enterprise (Private Cloud)	Service Providers (Public/Hybrid Cloud)
Local Backup 	<i>Example:</i> Local backup for multiple business units	<i>Example:</i> Hosted Applications including Backup as a Service
Replicated Backup 	<i>Example:</i> Remote offices with local backup	<i>Example:</i> Disaster Recovery as a Service
Remote Backup 	<i>Example:</i> Remote offices without local backup	<i>Example:</i> Backup as a Service over WAN

Figure 1: Cloud models enabled by secure multi-tenancy

WHY DATA DOMAIN FOR MULTIPLE TENANT ENVIRONMENTS?

Secure multi-tenancy provides logical isolation at the administrative and data path levels. This is the foundation of allowing the sharing of a Data Domain system and enabling access in a secure manner between different tenants on the same system.

DATA ISOLATION

Logical data isolation allows providers to spread the capital expenditure and operational expenditure of a protection storage infrastructure across multiple tenants. Data isolation is achieved by using separate DD Boost users for different MTrees or by using the access mechanisms of NFS, CIFS, and VTL. Figure 2 provides an overview of the secure multi-tenancy components.

ADMINISTRATIVE ISOLATION

A tenant-unit is a partition of a Data Domain system that serves as the unit of administrative isolation between tenants.

Multiple roles with different privilege levels combine to provide the administrative isolation on a multi-tenant Data Domain system. The main personas and roles are:

- **Data Domain Admin:** This role has full control to configure and monitor the entire Data Domain system.
- **Tenant User:** This is a new role for secure multi-tenancy and has privileges to monitor and report on only the assigned tenant units.
- **Tenant Admin:** This is a new role for secure multi-tenancy that is designed for monitoring, reporting, and limited modifications of the assigned tenant units. The Tenant Admin has more privileges than the tenant user.

The Tenant Admin and Tenant User can be restricted only to certain tenant-units on a Data Domain system and allowed to execute a subset of the commands that a Data Domain system Admin would be allowed. Both of these roles enable tenant self-service.

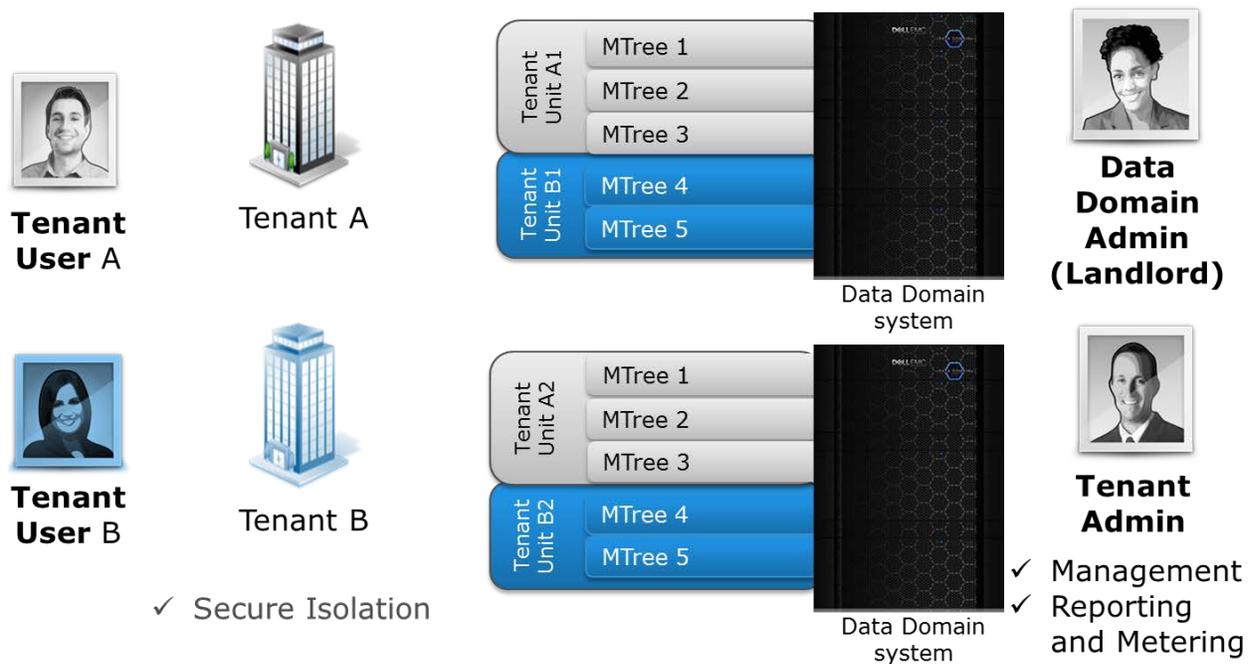


Figure 2: Components of secure multi-tenancy

For example, the Tenant User can monitor space usage, alerts, and performance for the privileged tenant units. The Tenant Admin is allowed additional privileges such as modifying the email addresses of the alert recipients and doing Data Domain fastcopy operations within the privileged tenant units.

In Figure 2, the Data Domain Admin has created tenant units for different tenants and assigned MTrees to those tenant units. The Data Domain Admin has also created and assigned Tenant Users and Tenant Admin for the tenant units in order to provide administrative isolation. Data access isolation is achieved by separate data access users for the individual MTrees.

BENEFITS OF ADMINISTRATIVE ISOLATION

Secure multi-tenancy provides the cost advantages of a consolidated resource with secure local tenant visibility & control. Secure administrative isolation allows the Data Domain Admin to control the infrastructure of the system and allow tenants to perform self-service to effectively monitor just the protection storage they have been assigned.

The hierarchy in the Data Domain Admin, Tenant Admin, and Tenant User roles allows for complex provider arrangements. For example, an infrastructure provider may want to be the Data Domain Admin, a reseller may want to be the Tenant Admin, and the end customer may want Tenant User access.

DATA PATH ISOLATION

The DD Boost protocol allows creation of multiple DD Boost users on a Data Domain system. With that, each tenant can be assigned one or more DD Boost user credentials that can be assigned access privileges to one or more MTrees in a tenant unit defined for a particular tenant. This allows secure access to different tenant datasets using their separate DD Boost credentials by restricting access and visibility. Data Domain uses SSL Certificate based authentication and data access is restricted through Tenant-unit's IP Addresses to tenant-unit's storage resources (Mtrees/Storage-Units).

Similarly, for other protocols such as CIFS, NFS, and VTL the native protocol level access control mechanisms can be used to provide isolation. For example, an active directory (AD) user or group can be assigned access to a CIFS share created on an MTree that is part of a tenant unit.

BENEFITS OF DATA PATH ISOLATION

In a shared environment, the provider must provide separate credentials for users to maintain the security of the datasets being protected. Data path isolation using access control mechanisms helps to achieve the goal of ensuring that a tenant gets self-service

access to only their dataset. The provider has the visibility of all the users accessing different tenant units and has the control of enabling or disabling access as needed.

MANAGEMENT

Tenant management relates to onboarding new customers including being able to provision data protection services in a timely manner and allowing tenant self-service.

One of the early challenges of onboarding new customers in a non-multi-tenant environment is the lead time of procuring, installing, and configuring protection storage in a datacenter. Secure multi-tenancy can shrink this lead-time from weeks to minutes by allowing the Data Domain Admin to do multi-tenant management on a Data Domain system.

Tenant self-service is important for large enterprises and service providers as they engage in IT as a service models with private, hybrid, and public cloud implementations. By empowering the end users, providers can gain the end customers' trust by enabling transparency and reduce the operational cost of providing services in a leveraged environment.

TENANT ONBOARDING AND PROVISIONING

Onboarding a new customer using secure multi-tenancy involves creation of a tenant unit on a Data Domain system. The Data Domain Admin can use the CLI to create one or more tenant units. A tenant may require multiple tenant units depending on the need for separating data path and monitoring access. For example, a tenant may want to align their tenant units with their organizational hierarchy; they can create a tenant unit for each business unit in their organization.

The Data Domain Admin can create and assign MTrees into the tenant units. One or more MTrees may be assigned into a tenant unit. Having multiple MTrees in a tenant unit allows different storage management policies for different workloads. For example, workloads in different MTrees may require different replication parameters for disaster recovery.

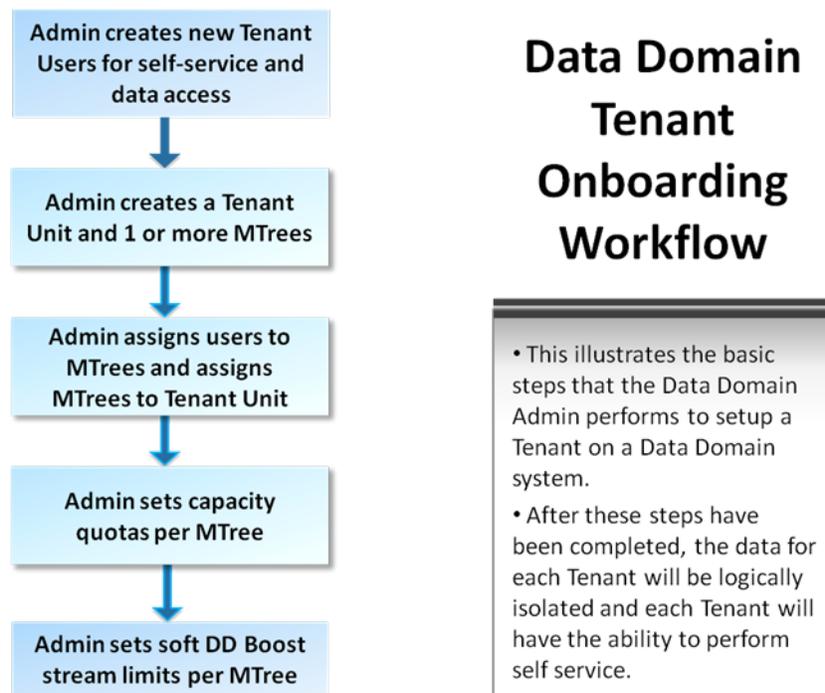


Figure 3: Tenant onboarding process

The Data Domain system allows for creating soft or hard capacity quotas for MTrees. This allows the Data Domain Admin to control the maximum capacity used by different tenants. Similarly, the Data Domain Admin can set soft quotas for DD Boost streams to control the number of streams that can be consumed on a per MTree basis. This allows the Data Domain Admin to control the maximum number of streams used by different tenants.

The Data Domain Admin can create new users for administrative isolation and assign privileges to the appropriate tenant unit.

BENEFITS OF TENANT ONBOARDING AND PROVISIONING

Secure multi-tenancy allows a tenant to span multiple Data Domain systems. The same tenant can own tenant units on different Data Domain systems. This allows the provider to use capacity on existing Data Domain systems before procuring new Data Domain systems and more importantly eases the planning process for onboarding new customers.

By leveraging capacity and stream quotas, the provider can easily manage Data Domain system resources and create different levels of data protection services in their service catalog.

Secure multi-tenancy can also shrink the data protection storage provisioning time from weeks to minutes which is a benefit for the provider and the user.

TENANT SELF-SERVICE

Using secure multi-tenancy, a Data Domain Admin can assign self-service privileges to different tenants. This allows a tenant to log in to the Data Domain system and only access the tenant units assigned to that tenant in order to monitor those resources.

A Tenant User has the privileges to monitor specific tenant units for important parameters such as space usage, streams performance, alerts, and status of replication context and snapshots. A Tenant Admin gets all the privileges of a Tenant User and additionally can modify the recipients list of alerts and also perform Data Domain fastcopy operations.

As an example, Figures 4 and 5 show how the Data Domain Admin can view logical capacity for all tenant units, and the Tenant User can only view the logical capacity of the assigned tenant units using tenant self-service.

```
sysadmin@dd860-15# smt tenant-unit show list
```

Name	Tenant Self-service	Number of Mtrees	Pre-Comp (GiB)
tu1	Enabled	2	3875.0
tu2	Enabled	1	1912.0

Figure 4: Data Domain Admin view of logical capacity reporting per tenant unit

```
tu1-usr1@dd860-15> smt tenant-unit show list
```

Name	Tenant Self-service	Number of Mtrees	Pre-Comp (GiB)
tu1	Enabled	2	3875.0

Figure 5: Tenant User view of logical capacity using tenant self-service

BENEFITS OF TENANT SELF-SERVICE

As large enterprises and service providers embrace the IT as a Service model, the dependence on a central support staff may become a bottleneck to service customer requests. When such requests are for information and management changes in the underlying

infrastructure, tenant self-service is particularly appealing to improve the time to service requests. Enabling such self-service in a multi-tenant environment provides benefits to both the provider and the consumer i.e. the tenant.

For example, when the provider bills the tenant on the basis of the amount of data protected, the customer can verify the historical usage over a period of time such as the last 30 days.

METERING AND REPORTING

Metering and reporting enable a provider to ensure they are running a sustainable business model. The need of such reporting in a multi-tenant environment is even greater for the provider to track usage on a shared asset such as a Data Domain system.

With secure multi-tenancy, the provider can get out of the box capabilities to track and monitor usage of the different resources of the system. Similarly the tenant can access metrics via tenant self-service. The tenant's view of the metrics is restricted to the resources that have been assigned to a particular Tenant User.

Different metrics can be extracted from the Data Domain system using SNMP. The SNMP MIB provides relationships of the different metrics to the tenant unit thereby allowing grouping the metrics on a per tenant basis.

Below are examples of metrics as viewed on the CLI by the Data Domain Admin and a tenant. They exemplify the restricted view of the Tenant Users.

LOGICAL CAPACITY REPORTING

Getting to know the capacity consumed on a per tenant basis is one of the most important metrics for a provider using a multi-tenant infrastructure. This information is used for billing customers. Data Domain systems can provide the pre-compressed logical capacity on a per MTree or per tenant basis. This is exemplified in Figures 4 and 6.

```
sysadmin@dd860-15# mtree list
```

Name	Pre-Comp (GiB)	Status	Tenant-Unit
/data/coll/backup	0.0	RW	-
/data/coll/lsu1	1415.0	RW	tu1
/data/coll/lsu2	0.0	RW	-
/data/coll/lsu7	1912.0	RW	tu2
/data/coll/lsu8	0.0	RW	-
/data/coll/lsu11	2460.0	RW	tu1

Figure 6: Data Domain Admin view of logical capacity reporting for all tenants

PHYSICAL CAPACITY MEASUREMENT

Physical capacity measurement measures the physical capacity consumed by a subset of files within the file system, based on how the files in the subset deduplicate with other files in the subset. This is a point in time measurement, based on when the measurement is requested.

You can specify the file system subset to measure in several ways: as a pathset (a set of files and directories), an MTree, a Tenant Unit (all files within a Tenant Unit), or a Tenant (all files within a Tenant).

The Data Domain system maintains a historical record of physical capacity measurements, which is available in tabular, graphical, and Excel spreadsheet formats, depending on the User Interface being used.

An example of a report in DD Management Center showing physical capacity measurement is shown in Figure 7 below.

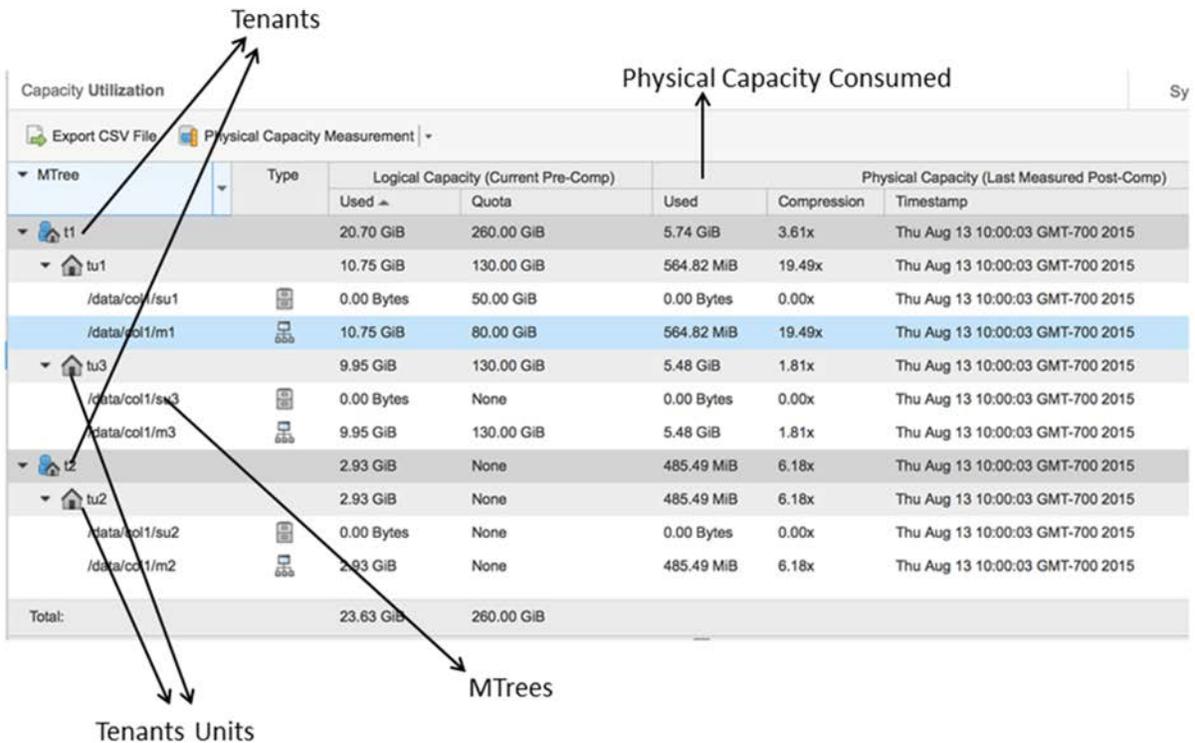


Figure 7: Physical capacity measurement

ALERT MONITORING PER TENANT

During daily operations, providers may want to monitor alerts on a per tenant basis in order to triage any issues being faced by specific tenants. Data Domain systems provide such alert monitoring for the Data Domain Admin as well as the Tenant User and Tenant Admin roles.

```

sysadmin@dd860-15# alerts show current
Id      Post Time          Severity  Class      Object
-----
19117*  Mon Mar  3 15:58:40 2014  WARNING  Filesystem Mtree=/data/coll/mtree1
19146   Mon Mar  3 16:44:50 2014  WARNING  Environment Enclosure=1:Sensor Id(0x)=51
19147   Mon Mar  3 16:44:50 2014  WARNING  Environment Enclosure=1:Sensor Id(0x)=55
19152   Mon Mar  3 16:50:08 2014  WARNING  Environment Enclosure=1:Sensor Id(0x)=61
19190*  Mon Mar  3 17:49:12 2014  CRITICAL Filesystem Mtree=/data/coll/mtree2
19201*  Mon Mar  3 18:07:49 2014  WARNING  Filesystem Mtree=/data/coll/lsu1
-----
(*) These alerts affect tenant-unit(s).
There are 6 active alerts.

```

Figure 8: Data Domain Admin view of current alerts on the system

```

tu1-usr1@dd860-15> alerts show current
Id      Post Time          Severity  Class      Object
-----
19117*  Mon Mar  3 15:58:40 2014  WARNING  Filesystem Mtree=/data/coll/mtree1
19201*  Mon Mar  3 18:07:49 2014  WARNING  Filesystem Mtree=/data/coll/lsu1
-----
(*) These alerts affect tenant-unit(s).
There are 2 active alerts.

```

Figure 9: Tenant user view of current alerts on the privileged objects only

HISTORICAL STREAMS PERFORMANCE PER TENANT

On a multi-tenant Data Domain system, different tenants may have different types of workloads and performance characteristics. In order for the Data Domain Admin to know such performance information on a per tenant basis, the Data Domain system can provide stream performance information on a per tenant basis for a defined period of time. Performance information is visible to the Data Domain Admin as well as Tenant Users via self-service.

```
sysadmin@dd860-15# mtree show performance tenant-unit tu2 last 1hours
INTERVAL: 10 mins
 "-" indicates that the data is not available for the intervals

Date          Time          Throughput          Streams
YYYY/MM/DD    HH:MM        read      write      rd/wr/r+/w+
              MB/s        MB/s
-----
2014/03/05 19:00      0.00     26.81      0/29/0/0
2014/03/05 19:10      0.00      0.00      0/0/0/0
2014/03/05 19:10      0.00     26.10      0/30/0/0
2014/03/05 19:20      0.00      0.00      0/0/0/0
2014/03/05 19:20      0.00     31.07      0/30/0/0
```

Figure 10: Historical stream performance per tenant unit

DD BOOST STREAM USAGE PER TENANT

The number of DD Boost streams on a Data Domain system is a limited resource and defines the capability of a tenant to backup or replicate to another Data Domain system. The Data Domain system allows reporting on DD Boost streams actively being used and historical DD Boost streams usage on a per tenant unit basis. This allows both the Admin as well as the Tenant Users to either triage performance issues or plan resource usage.

```
sysadmin@dd860-15# ddbost streams show active tenant-unit tu1
Name      Read      Write      Repl-out  Repl-in  Read      Write      Repl  Combined  Tenant-Unit
Streams  Streams  Streams  Streams  Limit    Limit    Limit  Limit
-----
lsu1      0         12         0         0         5         30        15    50        tu1
lsu11     0         0          0         0         -         -         -     -         tu1
-----
DD System Stream Limits: read=50 write=90 repl-in=90 repl-out=90 combined=90
sysadmin@dd860-15# █
```

Figure 11: DD Boost stream usage per tenant unit

FILE REPLICATION REPORTING PER TENANT

A multi-tenant Data Domain system can be used to support Disaster Recovery as a Service in a central data center for a large enterprise or a service provider. In such cases, it is important for the Data Domain Admin to know the replication status and usage on a per tenant basis. This is shown in Figure 12 below.

```

sysadmin@panda17# ddboost file-replication show active tenant-unit tul
File-replication Active for storage-unit: "lsul" tenant-unit: "tul",
03/10 15:07:13

Direction  Remote      Filename                                     Pre-comp   Pre-comp
           Host       -----                                     Size      Transferred
-----
Inbound    10.25.160.232 /data/coll/lsul/ostqa-ucsl1_1394488129_C2_F1_R1:1394488129:smt-wp:4:1:: 1,204,894,720 1,204,894,720
-----

                                     Inbound  Outbound
-----
Number of files                       1        0
Total pre-comp size (bytes)           1,204,894,720 0
Total pre-comp bytes transferred       1,204,894,720 0
Total pre-comp bytes remaining         0         0
-----
sysadmin@panda17#

```

Figure 12: File replication monitoring and reporting per tenant unit

BENEFITS OF MULTI-TENANT REPORTING METRICS

The per-tenant metrics on a multi-tenant Data Domain system provide several benefits to providers:

- Ensures the appropriate usage of their protection storage assets based on capacity and performance
- Provides a basis for front end logical capacity or physical capacity used based chargeback and billing to tenants
- Provides alerts in a global and per tenant fashion to view the inter-dependencies of such alerts and do incident triaging on a per tenant basis
- Enables more efficient capacity & migration planning and assignment of the appropriate resources while onboarding new customers

The capability of tenants to view metrics for the resources they are assigned provides them with similar benefits. This also allows for a reseller model in which the reseller is providing monitoring services as a value add in their service portfolio.

CONCLUSION

After reading this paper you should have a better understanding how the secure multi-tenancy for Data Domain feature can dramatically improve protection storage efficiencies for large enterprises and service providers.

To summarize, secure multi-tenancy will help you:

- Achieve secure data isolation by tenant on shared Data Domain systems
- Reduce data protection storage costs
- Shorten time for service requests
- Enable efficient IT as a service for Data Domain protection storage
- Provide tenant self-service administration and reporting
- Give the provider control over capacity and stream count resources assigned to each tenant on shared Data Domain systems
- Perform chargeback/billing using logical front end capacity or physical capacity used by tenant
- Deliver protection storage as a service in a private cloud environment and for service providers in a hybrid or public cloud environment
- Measure and report on physical capacity used by tenant for effective capacity planning and migration planning

If you would like to know more about Data Domain systems, refer to our [Data Domain Data Invulnerability Architecture](#), [Data Domain SISL](#), [Data Domain Replicator](#), [The Business Value of Data Domain](#), [The Business Value of Data Domain Boost](#), [The Business Value of Data Domain for Oracle](#), and [Data Domain Boost for Oracle RMAN](#) white papers found on EMC.com. Also, if you have an iPad, check out the [Data Domain Interactive Product Guide](#) on EMC.com. Finally, please join us on [The Protection Continuum blog](#) to discuss this and other EMC data protection & availability topics. You can also visit the [Dell EMC Store](#) to explore Data Domain products.