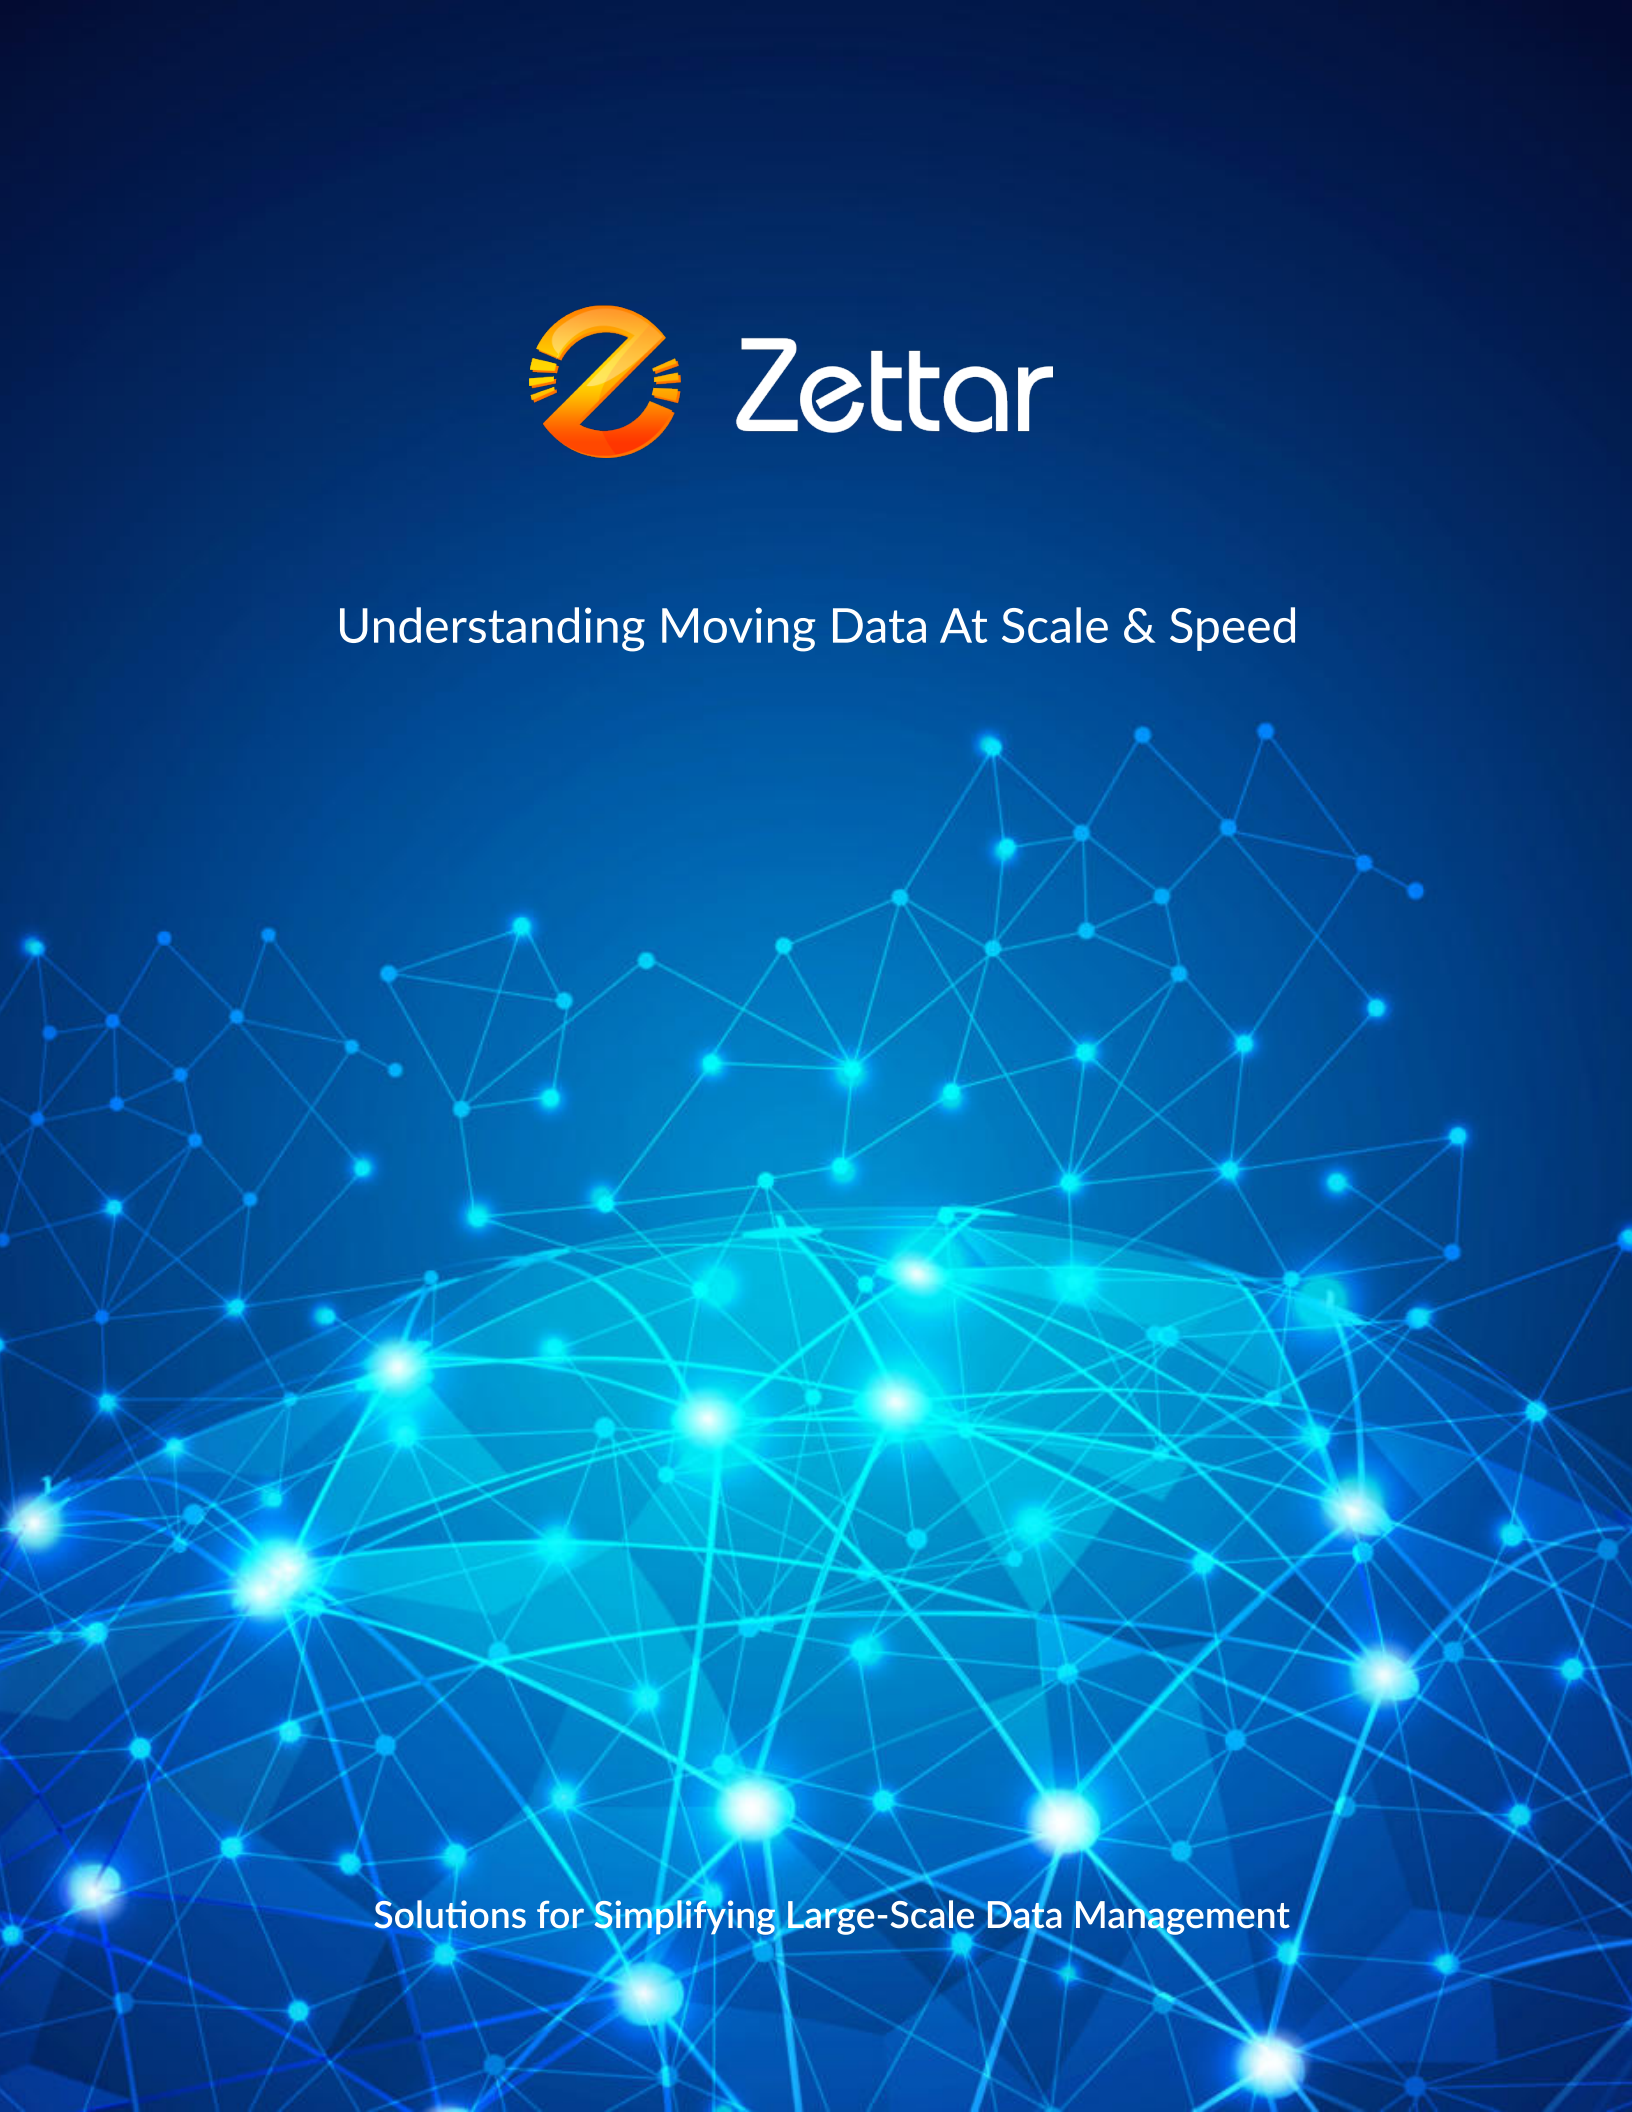
Zettar
Understanding Moving Data At Scale & Speed
Solutions for Simplifying Large-Scale Data Management

Zettar

Contents

Introduction

Zettar was founded on the idea that the world faces exponential data growth, mostly due to the massive amounts of machine generated data. To address such rapid data growth takes a different thinking and approach. In general, most people would regard moving data as something trivial. Some may view it as a software alone or network alone task. Such common misconceptions have been further encouraged by a few commercial data mover software vendors who claim that their legacy Managed File Transfer software, together with their proprietary data transfer protocols, would be able to address any data movement challenges at maximum speed.

In today's hyper-competitive markets, organizations need flexible infrastructure, workflows are becoming increasingly complex, and data sets are continuing to grow unchecked. Those legacy incumbent data mover software solutions force enterprises to architect overly complicated and costly systems that reduce IT agility. As a result, important business insights remain locked away, out of reach of decision makers.

Zettar has introduced a set of modern data mover software products, all integrated into a single application, namely `zx`, to address the formidable data growth challenges faced by all modern distributed data-intensive businesses. This white paper first presents an analogy that orients the readers. Then, it offers three actual examples to further enhance understanding.

Understanding moving data at scale & speed

Set up the infrastructure correctly

Foremost, it's imperative to keep in mind that moving data at scale and speed requires an infrastructure that is correctly set up, all stacks properly tuned, and operating well with the confidence that it has the capacity, reliability, and performance level sufficient for the desired data rates. Software is at most 1/4th of the equation. This can be well illustrated with the following.

A Water Transport Analogy

Thus, unless there is a sufficient understanding of the environment in which the software is to run, it's impossible to achieve a good outcome.

This is the reason why we advise you not to trust typical vendor offered "tools" such as file transfer calculator, numbers based on benchmarks conducted in a lab, A/B testing results etc. If estimates are not generated in your own environment, their artificial values are not of much use to you. If you really wish to see how a data mover performs, evaluating the software in your own environment is the most reliable way. The most essential concept to grasp is the co-design principle: integrated consideration of storage, computing, networking (including network security), and the scale-out and concurrency properties of the software data mover application.

As the water transport analogy shown above, the most important element is to have sufficient storage throughput to sustain the desired data rate(s). A dry reservoir will not enable a water transport system to create good water flows, regardless how well the three other elements are functioning.

The endeavor is not a software or network alone task

Furthermore, some commercial vendors over the years have spread the myth that UDP is superior than TCP. This is patently false and completely contradicts the reality. UDP is only defined in the Internet Engineering Task Force (IETF) Request for Comments (RFC) 768. Once anything else is introduced, especially the typical mechanisms built-in TCP such as windowed flow control, congestion avoidance, and slow-start, the result is at best a "home-grown TCP". It is no longer UDP anymore. Any good network programming books, graduate level network programming courses taught in good universities, or online courses should make this aspect clear. Zettar builds its application level transfer protocol on stock modern TCP, which has evolved very well for moving data at scale and speed under many circumstances, as the few examples demonstrate later.

Again, referring to the water transport analogy above, a complete data movement solution always involve storage, computing, networking, and data mover software. By itself, software acts just like water pumps in a water transport system. Regardless how well they are engineered, by themselves water pumps cannot move water. The same is true for transporting data.

Nevertheless, just like quality water pumps help generate good water flows, well engineered software data mover applications help use the underlying infrastructure elements efficiently, thus accelerating workflows.

The three different categories of data movement solutions

Confusion about the nature of moving data in the modern times have caused many other misconceptions. We will first briefly describe the three categories, then a few common misconceptions and the correct thinking.

The three data movement categories:

1. **Data access solutions** are used mostly by end users, who as a rule work with data interactively. The involved data sizes tend to be small. The majority of the data stays where it is.

2. **Data transport solutions** in modern times are the most important category and is the foundation of the 3rd category. Anytime data is generated at a fast rate and high volume, yet local storage capacity or processing resources are insufficient to handle the data, then moving it efficiently and timely to another location becomes critical. This is where data transport solutions come into the play. With most data machine generated these days, the above scenario occurs far more frequently than most realize.

A simple to visualize example is a large telescope, which usually is located on a remote mountain top. Geographically and cost-wise it's simply not appropriate to set up local storage and processing facilities in the limited space available at the mountain top. As a result, the large amounts of super-high resolution camera images must be transported to a data center or cloud for further processing.

3. **Data traffic aggregation solutions** are a well-practiced network traffic routing technique. By aggregating the traffic flows generated with different data transport solutions and routing such traffic over a few well-defined network connections, the utilization of such connections can be boosted to an outstanding level (often > 90% and Tbps). One can see it as the reverse of the classical "divide and conquer".

A few more common misconceptions

- **Create data lakes.** By the incorrect impression that the end-user is the most important element in today's data utilization scene, the term "data lake" was coined and is used more and more. There are serious economic reasons that this is a faulty thinking. Storing data in huge volumes to form a "lake" implies a large deployment of storage systems or high cost if stored in a public storage cloud. In the former case, other than the high operations costs, when time comes to do technology refresh (*3 to 5 years for an on premises deployment*) it's going to be a tedious and costly IT pain.

  Furthermore, by the 20/80 percent rule, massive amounts of data demand copious computing power, which is expensive to acquire, to setup, to operate, and to support. In most cases, a far better option is to keep data mobility high to avoid accumulating so much as to form a "lake" and store massive amounts of data only if the situation warrants. For example, is it reasonable for a university lab to create a data lake? Absolutely not! But is it reasonable for a large and established national supercomputing center to create such a "lake"? Definitely! Use the appropriate approach for a given situation is the key.

- **Bringing computing power to data.** Again, this is a biased view likely caused by incomplete understanding of the nature of data growth in modern times. It is still uncommon for data to be generated right where the processing power is in proximity. For example, even in the cloud, some large-scale simulations may generate massive amounts of data, which still needs to be stored in the cloud provider's cloud storage. Nevertheless, the storage is not really "right where the processing power" is. Thus, having a proper data mobility solution is still essential to your IT project even it's done in the cloud.

- **An unreasonable fascination with virtual machines and containers.** Neither is appropriate for heavy I/O driven workloads such as moving data at scale and speed. A rule of thumb: anytime the desired data transfer rate is ≥10Gbps, use physical servers dedicated to data movement whenever feasible. Note that the several leading public cloud service providers: AWS, Microsoft Azure, and IBM offer physical instances. Why? Other than some of their customers need such instance types, the providers need them for their internal projects too.

Real world product trial and competition results

Here you will find a few recent world leading engineering accomplishments achieved using Zettar zx. They are excellent examples of what are feasible in real world production environments.

1PB in 29 hours over 5000-miles; 94% bandwidth utilization

A few notes first: the Energy Sciences Network (ESnet) is a production network operated by the U.S. DOE, so for the trial there was a 80Gbps hard bandwidth cap. Furthermore, full TLS encryption is used, together with checksumming. The data transfer setup is very modest, please see the photo below:

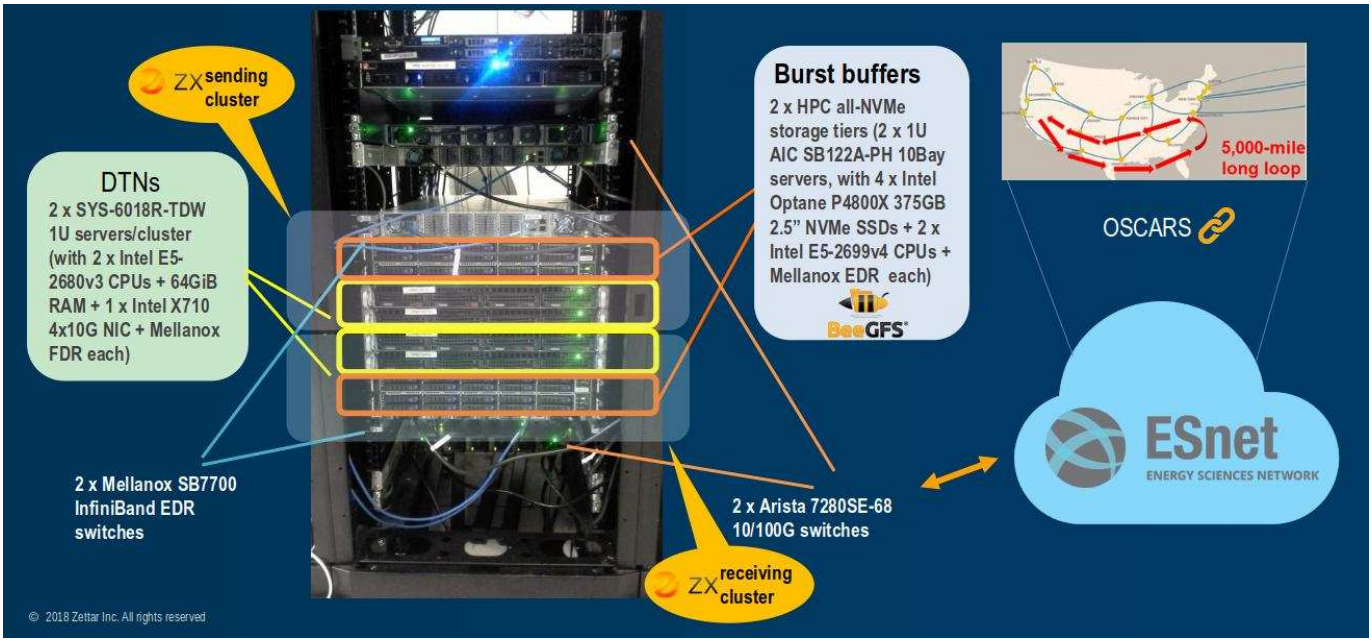

It's also of interest to note the flat transfer profile the pre-production trial generated. Please see below - from the ESnet's Network Overview Portal, which is world-wide visible by the public.

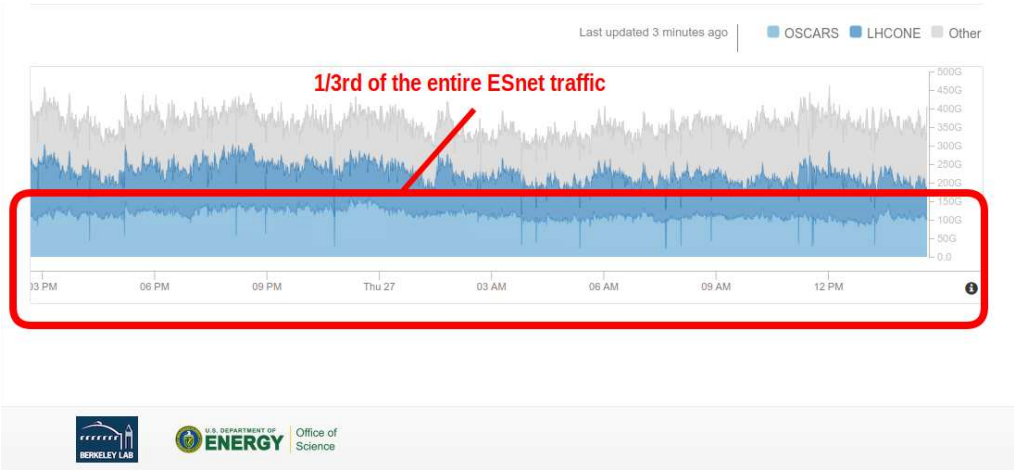

Winning the Supercomputing Asia 2019 (SCA19) Data Mover Challenge (DMC)

Please note that the fiercely competitive SCA19 DMC is the only international competition at the highest level of this kind in the most recent decade. Zettar competed against six other elite national teams from the U.S. and Japan, including the two leading DOE national labs Argonne National Lab and Fermilab, together with Japan's JAXA + Fujitsu team. Quoting the SCA19 DMC Press Release, https://bit.ly/3C0O4W2, "*Overall Winner Award: Zettar Inc. They achieve an amazing 68 Gbps transfer rate between Chicago and NSCC (Singapore).*" Note that this feat was accomplished with a very restrictive environment, so proper tuning was not feasible. Also see below for the extremely demanding network topology.

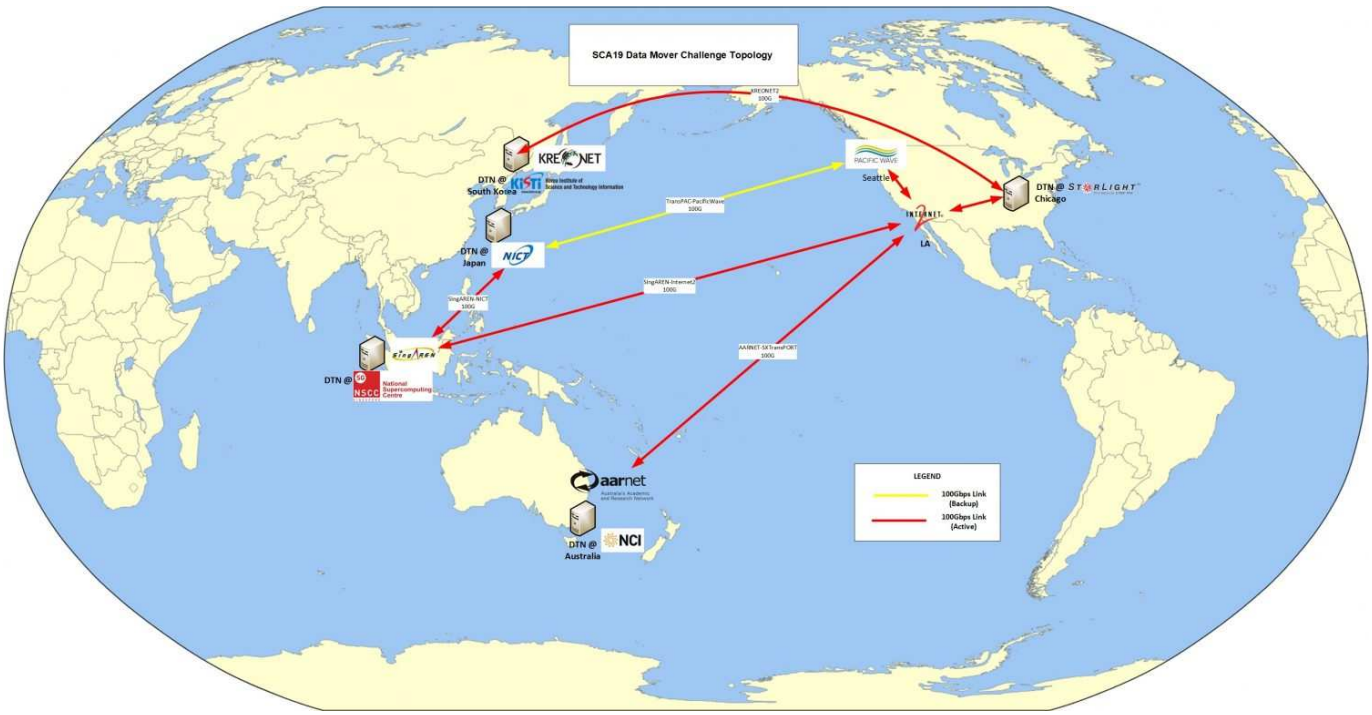

A Historical 1st data transfer production trial between Poland and Singapore

In early November 2019, barely two weeks prior to SC19, Zettar zx was used to established the historical 1st Poland to Singapore data transfer production trial over the new CAE-1 100Gbps connection, which links London and Singapore across a vast distance of 12,375 miles, https://bit.ly/2W7NK4A. The trial achieved a world-leading average transfer rate of ~60 Gbps using modest production set up and hardware, including two production Lustre parallel file systems. Both were built using only conventional HDDs. The trial used a single data transfer server at each end. The outcome was shown in the ICM SC19 booth 1393 and the A*STAR NSCC Singapore booth 2043. The trial was also presented in the Supermicro's SC19 booth as an invited talk.

Key takeaways

There are three essential key takeaways from the information presented above:

1. **The co-design principle is effective.** It means integrated consideration of storage, computing, networking (including network security), and concurrent, scale-out data mover software, is a main reason why those accomplishments were attained. None of them is a network alone or software alone exercise. Zettar applies the principle whenever it can.

2. **Proprietary data transfer protocol is unnecessary.** Zettar `zx` does have its application level data transfer protocol, but the basis is TCP. In fact, modern stock TCP is so good that when properly used, network latency is not an issue.

3. **Scale-out data mover enables a cost-effective and forward-looking setup.** Reviewing the 1st accomplishment and the hardware setup employed, it is evident that most hardware components were not of heroic class, brand new high-end types. The use of a scale-out capable data mover is a key reason how it was made happen. Note that in 2018, Intel DCG EVP Navin Shenoy stated that the world faces exponential data-growth and 90% of the world's data was generated in the past two years (i.e. since 2016). Yet the majority of the data mover software applications available today are still of the traditional host-oriented design, often more than two decades old. None of them are capable of providing such a ROI benefit to your existing IT hardware investments. It is time to leverage a scale-out data mover that can tackle the exponential data growth now and in the future.

Zettar

Summary

Zettar zx is a petascale-proven, scale-out, hybrid-cloud native software data mover solution. It adapts to your environment and application needs. The solution is ideal for the performance and scale needs of Life sciences, Oil & Gas, HPC, large scale Network Attached Storage (NAS) tech refresh and/or data migration, autonomous vehicle data collection, AI, Machine Learning, Media & Entertainment post-production workflow acceleration, or other large bandwidth and I/O intensive applications that require moving data at scale and speed. zx brings in flexible and superior data mobility to any large-scale data management endeavors, requiring much less operational overheads and fewer hardware resources compared to incumbent solutions.

Zettar zx addresses the acute data gravity problems by providing a fast, efficient, and resilient distributed data mover system that is hybrid cloud native and delivers simplicity, scalability, and world-leading performance at the same time. Part of zx's ease of use in a hybrid cloud environment includes rapid provisioning to reduce time to get new workloads deployed, along with elasticity scaling, resiliency, performance, and cost effectiveness.

We invite you to visit www.zetttar.com to learn more or request a trial or demo.