

Beyond Legacy Tape: Expecting More from Mainframe Virtual Tape

Art Tolsma
Luminex

November 2019
Session **DA**



Beyond: Expect More

- Do More. Expect More. Embrace More.
- Enterprise IT is being fundamentally transformed
 - Cloud computing is Driven by Agility and “Do More” and not “For Less”
- The IBM Mainframe remains critical on-premise Enterprise IT
- “It’s All about Integration” and “Significant focus to make the mainframe a seamless participant” – Jeff Magdall SHARE Phoenix Session March 2019

Luminex At A Glance

Delivering mainframe data solutions worldwide for 25+ years

- Enabling mainframes to leverage distributed systems resources via native I/O channels
- Renowned for its industry-leading, enterprise-class mainframe virtual tape solutions

MISSION STATEMENT

Luminex serves as a trusted advocate helping *enterprise customers* **protect**, **manage**, and **leverage** corporate data assets by developing and delivering high quality, innovative technology solutions.




What are Luminex Customers Doing with MVT?

CAPACITY

**5.7
PB**

of host-view
tape capacity for a
single subsystem



OF VOLUMES

**12
Million**


in a single tape
management catalog



THROUGHPUT

**10.8
GB/s**


in a single site's
tape subsystem



TAPE MANAGEMENT CATALOGS


**39
TMCs**

in a single installation



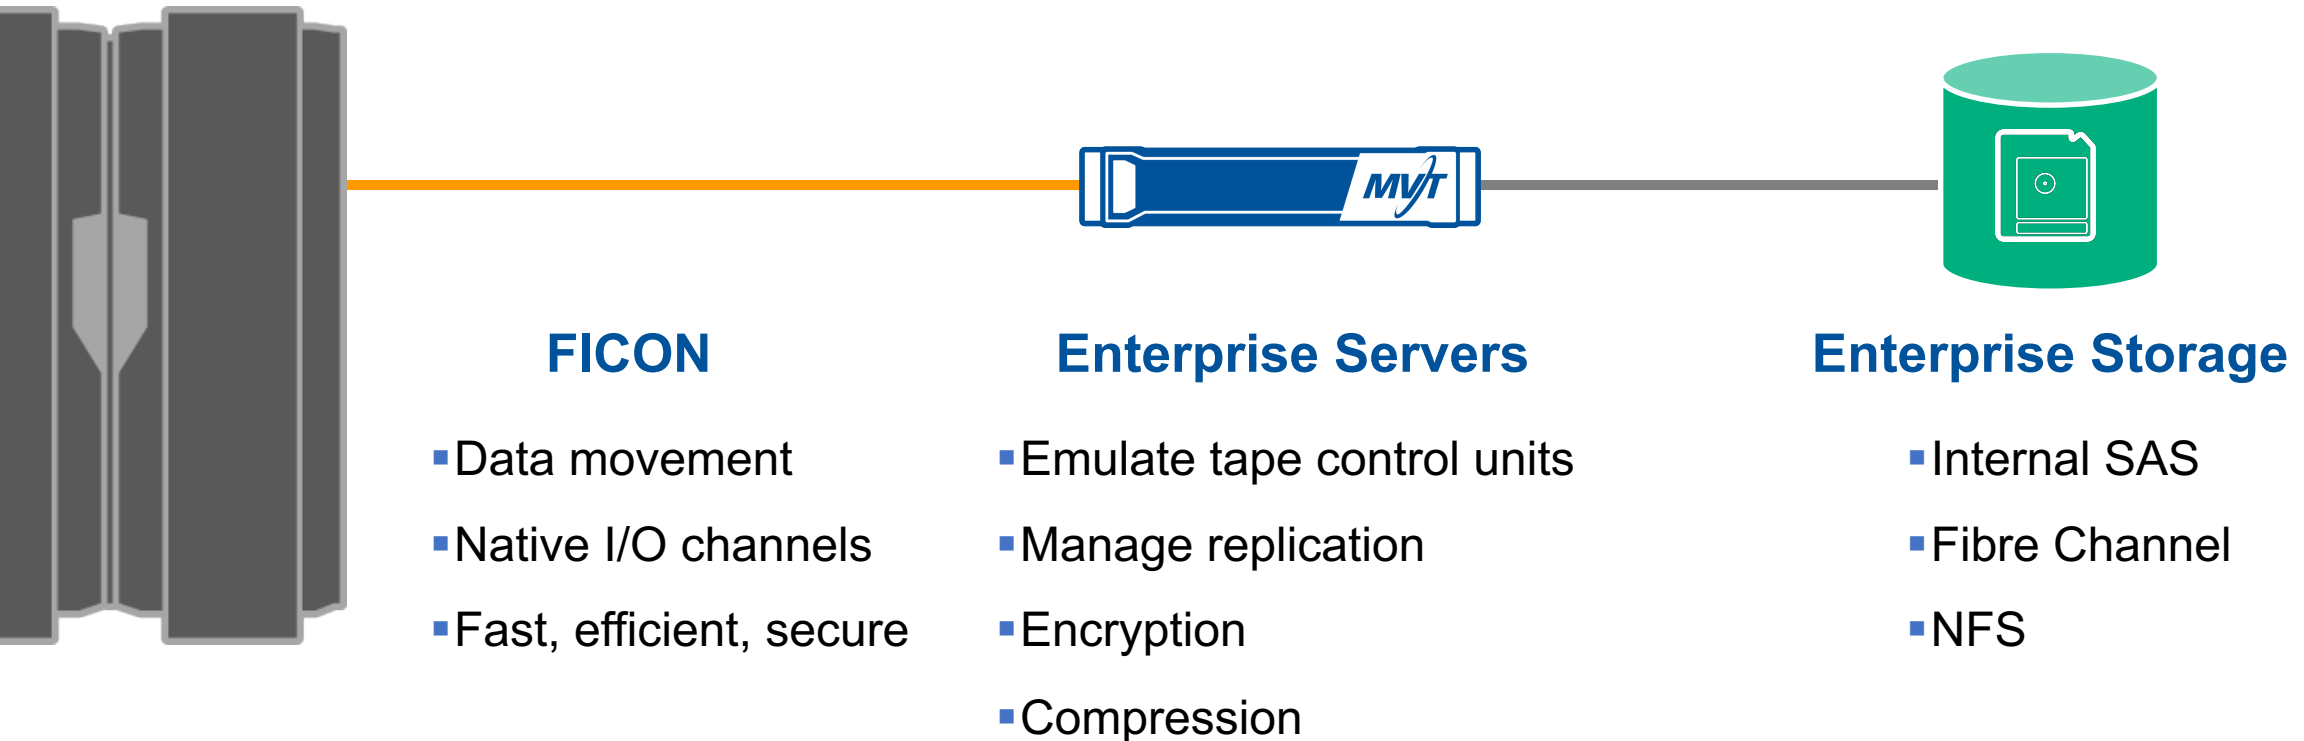
Traditional Tape Use Cases

- Primary operational
 - Batch workloads
- Backup and recovery
 - Ship tapes to an off-site warehouse
 - Ship from warehouse back to production or a DR site
- Archive & compliance
 - Ship tapes to an off-site warehouse
- Data sharing
 - Cut a tape and ship it to business partners



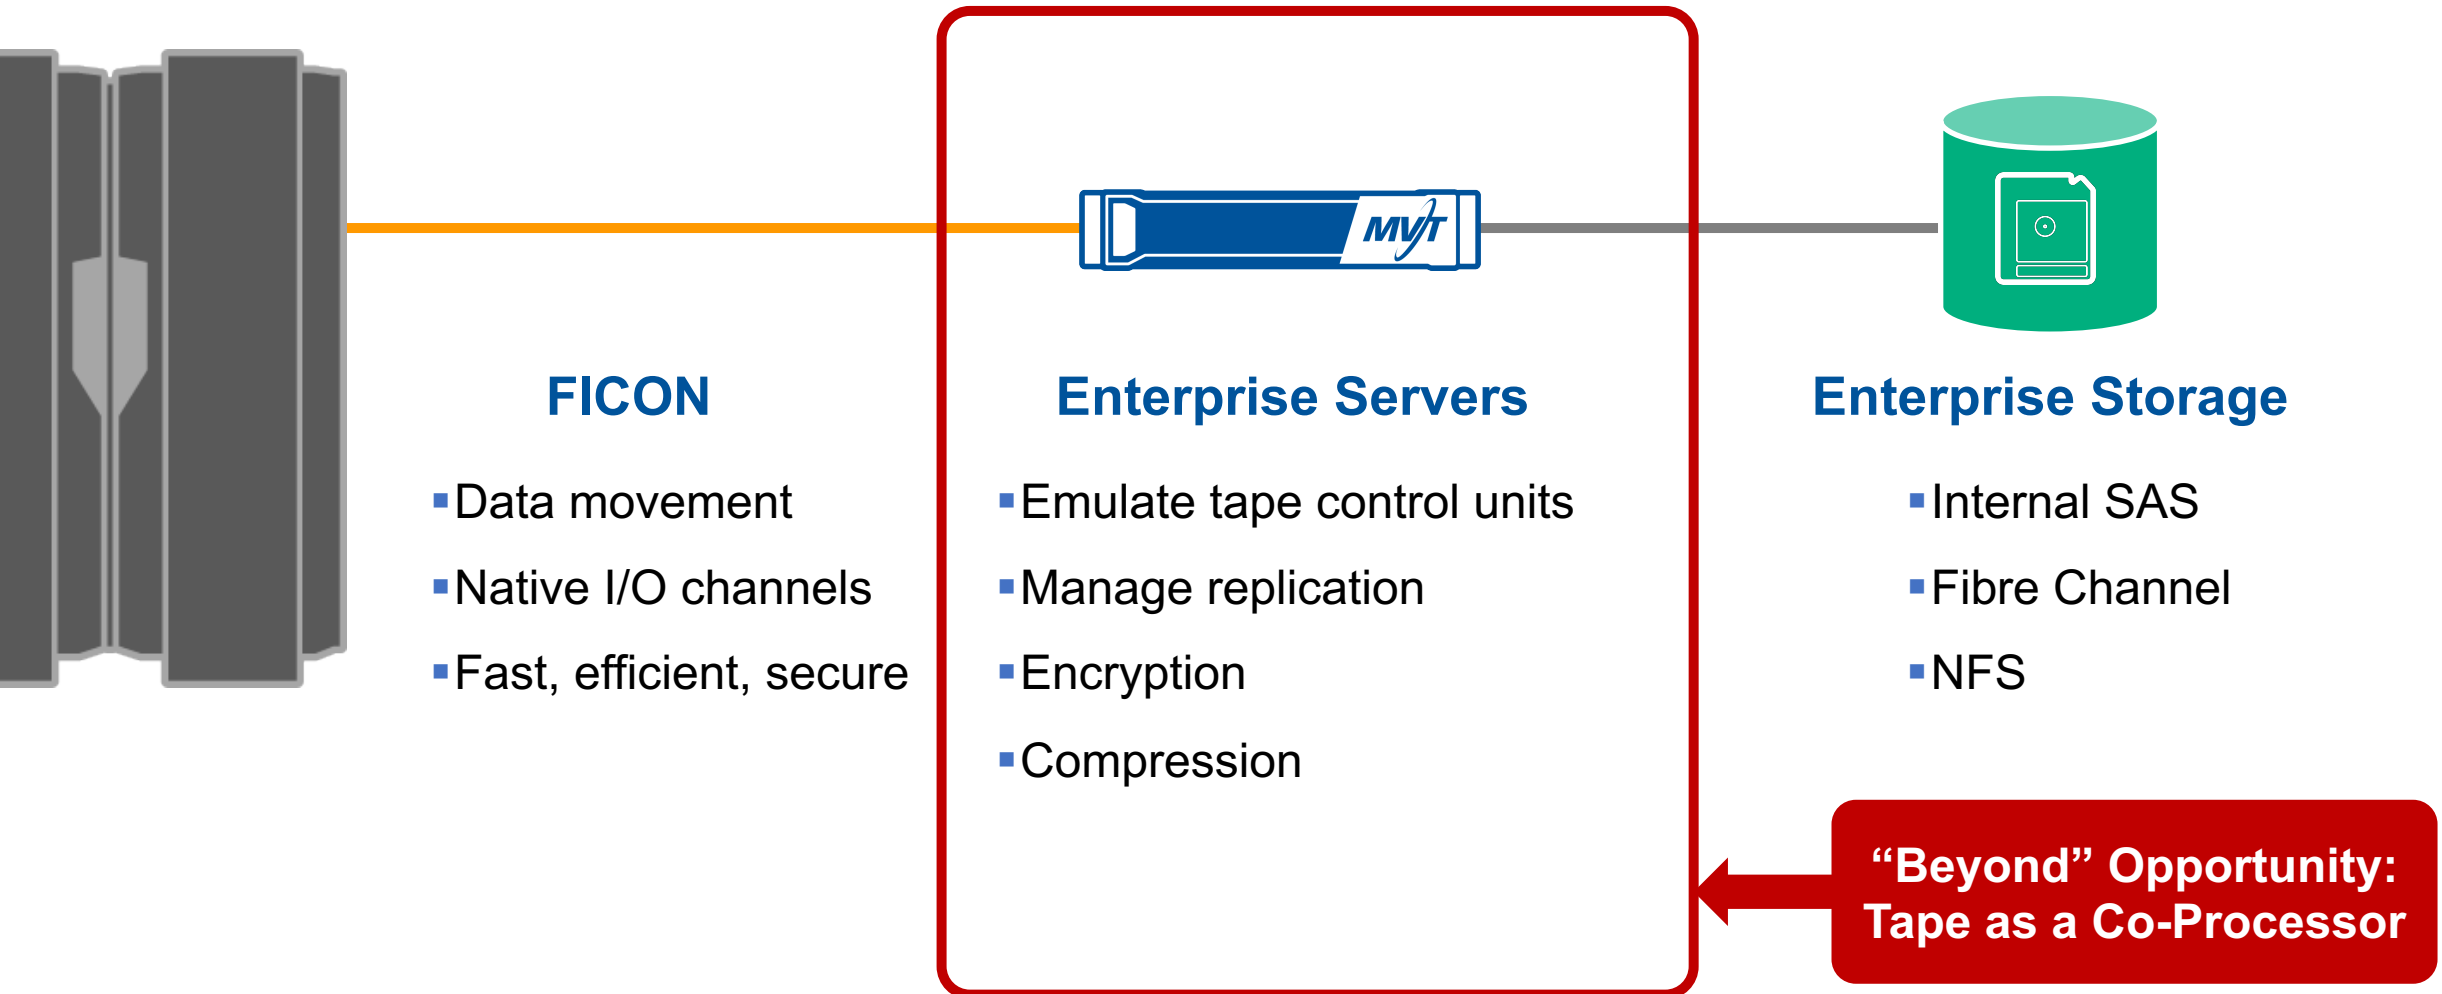
Components of a Virtual Tape System

Most major tape solutions use commodity hardware:



Components of a Virtual Tape System

Most major tape solutions use commodity hardware:



-




Push Button DR

RepMon: Replication Monitor

Provides **real-time status monitoring and logging** of virtual tape data writes and replication to a remote location at the VOLSER level

- Identifies write and replication status of mainframe tape VOLSERs
- Identifies if off-site virtual tape data is still consistent with the primary datacenter
- Provides visual and audit capabilities to confirm when tapes are successfully protected off-site
- CloudTAPE Dashboard adds versioning and tiering status of VOLSERs in the Cloud



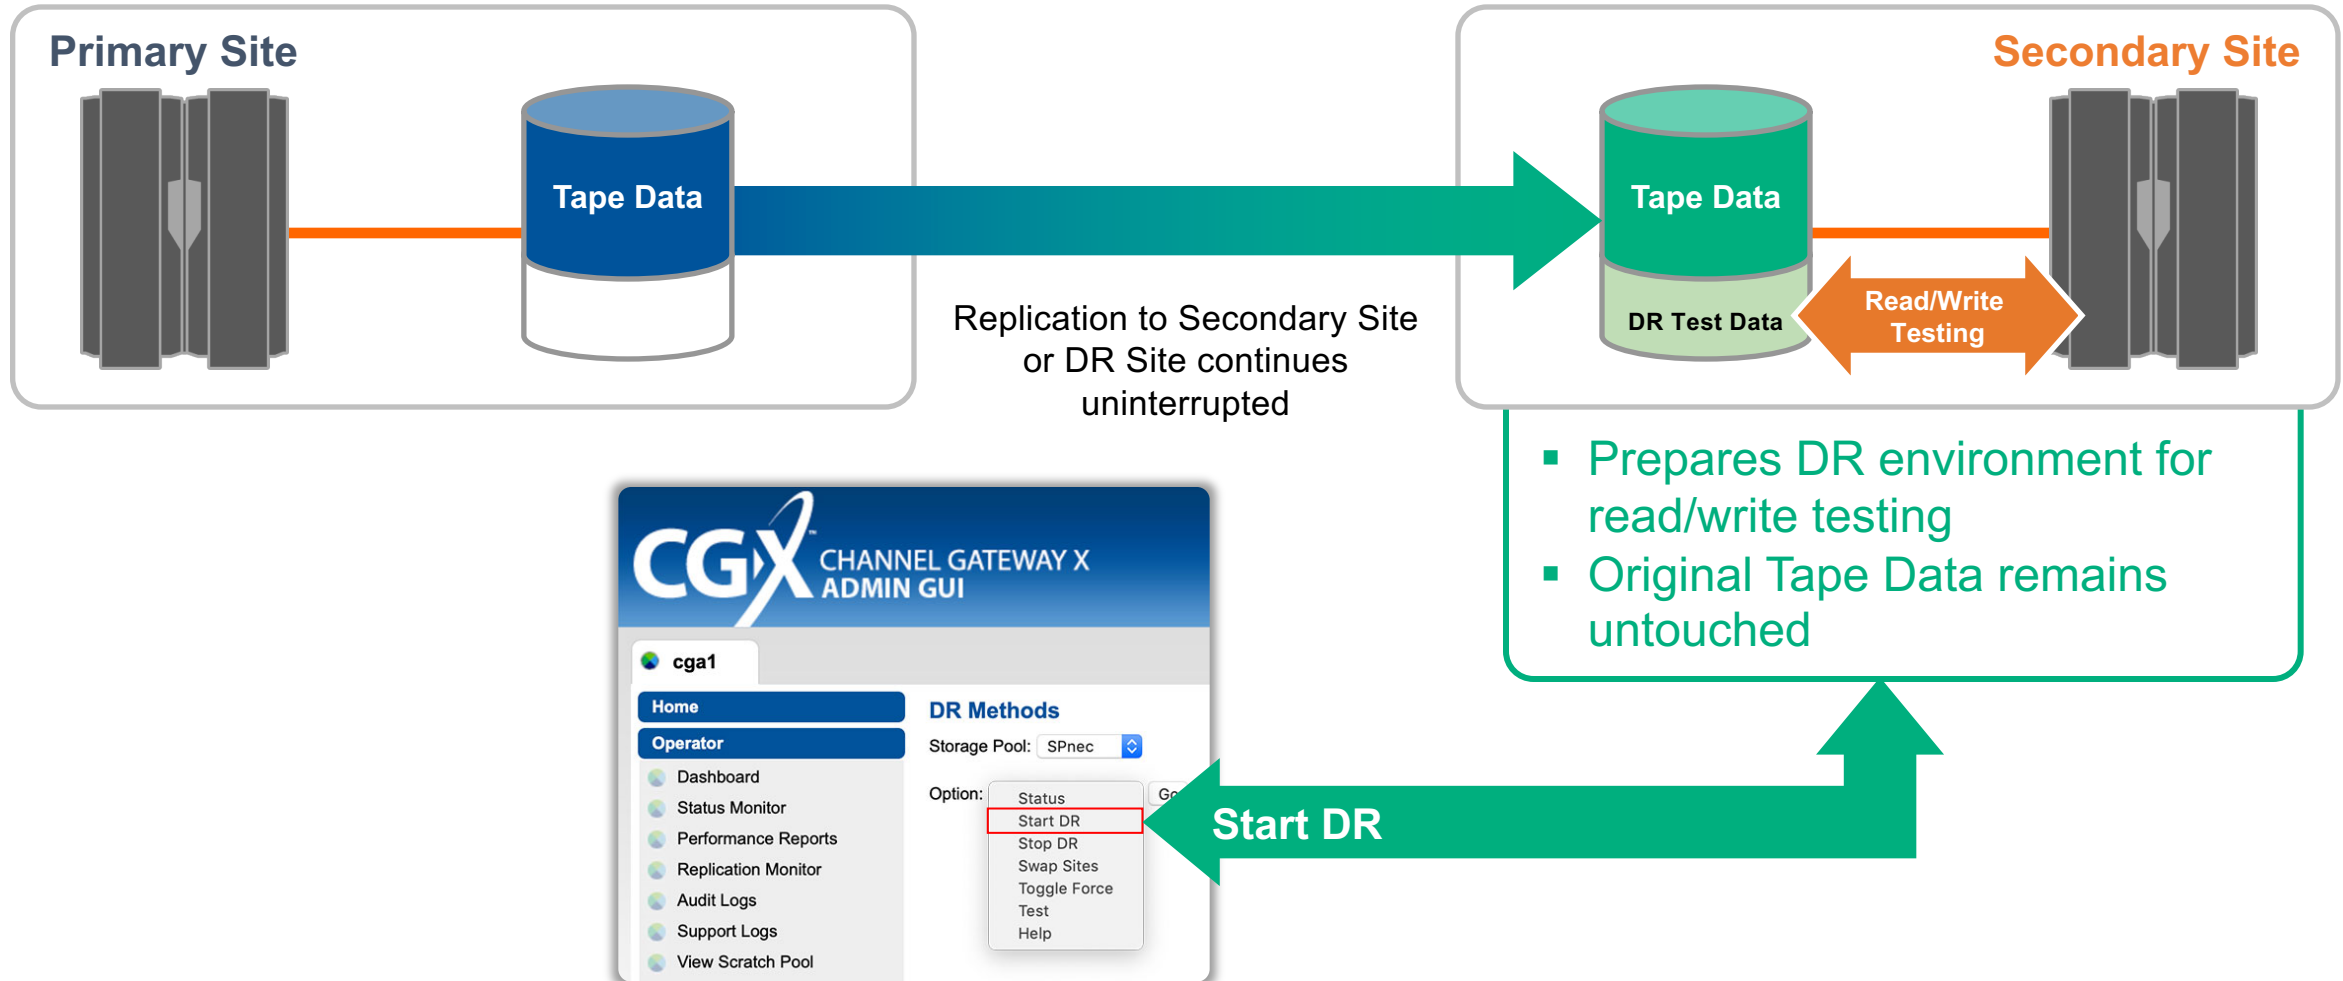
Push Button DR Testing

Replication During Normal Operations



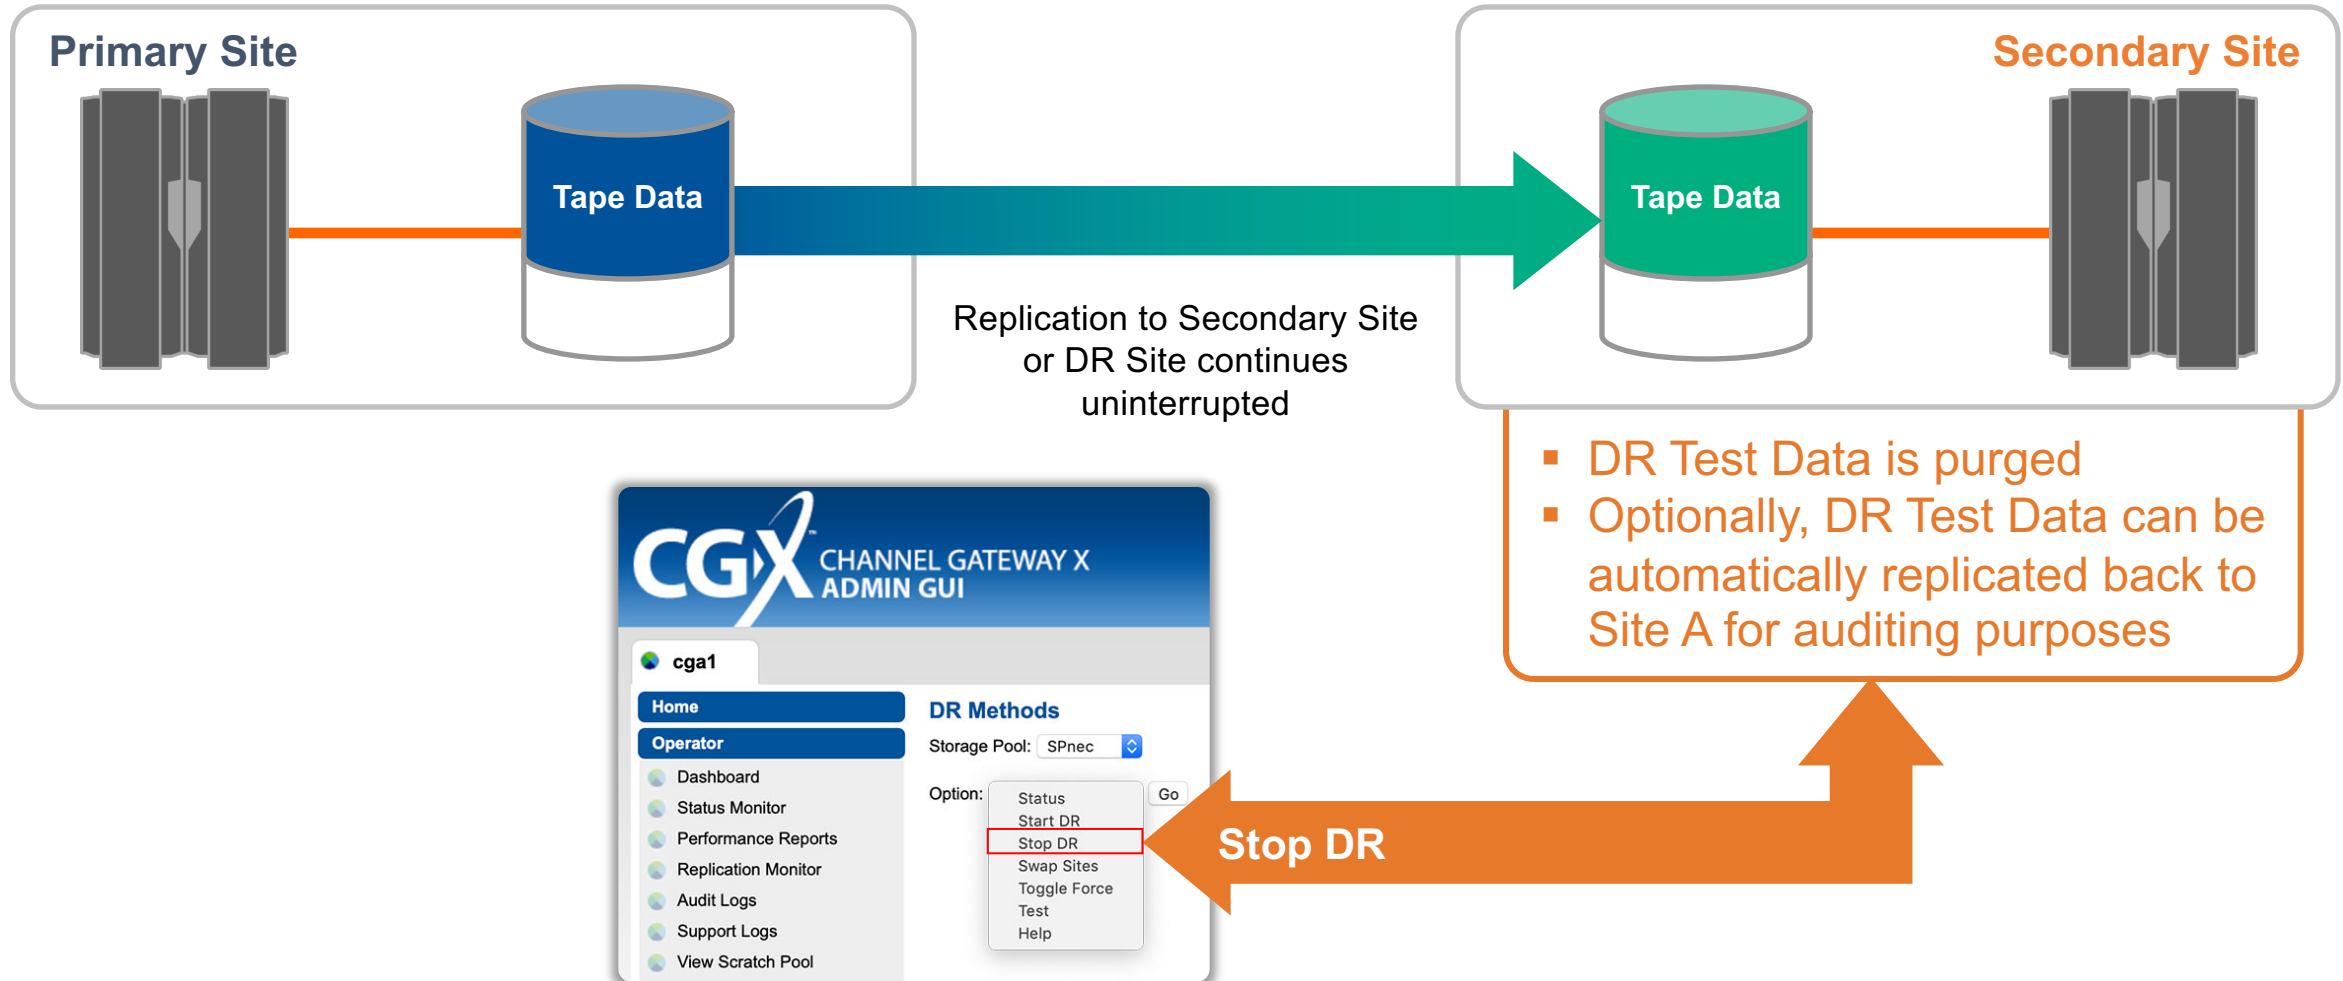
Push Button DR Testing

Replication During DR Testing

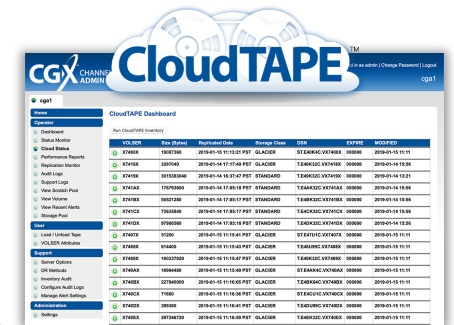



Push Button DR Testing

After DR Testing is Completed



- Provides an **air gap** for virtual tape
 - Protect tape data with programmatic immutability
 - Support for cloud providers' versioning capabilities with reporting
- Reads/writes work the same from the mainframe perspective
- Support for cloud and object storage providers' standard APIs
 - AWS, Azure, Hitachi HCP, NetApp StorageGRID, etc.
- Transparent tape data movement
 - No additional mainframe resources required (encryption, movement & management)
 - Seamless support for cloud storage tiering / recalls (e.g. S3-to-Glacier)
- DataStream Intelligence support
 - Tape metadata added to the stored object's metadata (creation/replication dates, tape size, tape label)
 - Metadata can be leveraged by cloud-based applications
- Luminex GUI tracks and reports tape versions, status and metadata



Multitenant Views in Admin GUI

- Separate **Global** and **Storage Pool Level** views
- Storage Pool administrator can only see information for Storage Pools and devices assigned to them
- Global administrator can see all information for all Storage Pools
- Global administrator can assign capacity “quotas” for individual Storage Pools
 - These quotas are reported and alerted on but are not strictly enforced
 - E.g. a tenant can use 120% of their quota as long as there is enough backend storage

Storage Pool Level View

Restricts Display to Assigned Storage Pool

- Status Monitor for Associated Devices
- Replication Monitor
- Replication Audit Logs
- View Scratch Pools
- View/Search Volumes
- Capacity Information
- Compression Rates
- Load/Unload Tape
- VOLSER methods
- DR Methods
- Inventory Audit
- Admin Settings

Global View

Display Extended to All Storage Pools

- All of the Storage Pool Level view functionality
- See all of the **storage pools** configured on the system
- **Execute operational commands**, such as restarting the MVT
- See **historical and real-time global performance statistics**, such as I/O or network rates
- Generate and download **support logs**
- View and modify **alert settings**
 - This does not imply that alert thresholds cannot be distinct among tenants, only that the management of alerts is done by the Global administrator

Multitenant GUI – Global View

The screenshot displays the CGX Channel Gateway X Admin GUI. The top header bar is blue and contains the CGX logo, the text "CHANNEL GATEWAY X ADMIN GUI", and user information: "Logged in as admin | Change Password | Logout" and "cga1".

On the left side, there is a sidebar with a "cga1" tab. Below the tab are two main sections: "Home" and "Operator". The "Operator" section contains a list of menu items, each with a small icon: Dashboard, Status Monitor, Cloud Status, Performance Reports, Replication Monitor, Audit Logs, Support Logs, View Scratch Pool, View Volume, View Recent Alerts, and Storage Pool (which is highlighted).

The main content area is titled "Storage Pool" and contains a table with the following data:

Storage Pool	Total Capacity	Total Used	Total Available	Percent Used	Comp Ratio
CustomerA	400.00TB	248.40TB	151.60TB	62.1	3.26
CustomerB	200.00TB	147.20TB	52.80TB	73.6	4.16
CustomerC	150.00TB	88.95TB	61.05TB	59.3	3.31
CustomerD	50.00TB	34.45TB	15.55TB	68.9	3.64


Multitenant GUI – Storage Pool Level View

The screenshot displays the CGX Channel Gateway X Admin GUI. The top header bar is blue and contains the CGX logo, the text "CHANNEL GATEWAY X ADMIN GUI", and user information: "Logged in as admin | Change Password | Logout" and "cga1". Below the header, a sidebar on the left lists navigation options: Home, Operator, Dashboard, Status Monitor, Cloud Status, Performance Reports, Replication Monitor, Audit Logs, Support Logs, View Scratch Pool, View Volume, View Recent Alerts, and Storage Pool (which is highlighted). The main content area is titled "Storage Pool" and contains a table with the following data:

Storage Pool	Total Capacity	Total Used	Total Available	Percent Used	Comp Ratio
CustomerA	400.00TB	248.40TB	151.60TB	62.1	3.26


Synchronous Tape Matrix (STM)

- Continuous Availability
 - Resilient architecture instantly and automatically adjusts to multiple failures without interruption
 - Data is always available for I/O
 - No downtime from failover or restore processes
- No mainframe components or processes to manage
- No idle components to buy
 - All components contribute to day-to-day operations, not just during failure events
- Easy to implement
 - No host scripts or policies required
- Scalable
 - No limitations for throughput, capacity or degrees of redundancy
- Modular design ensures investment protection
- Supports heterogeneous storage systems and data reduction technologies

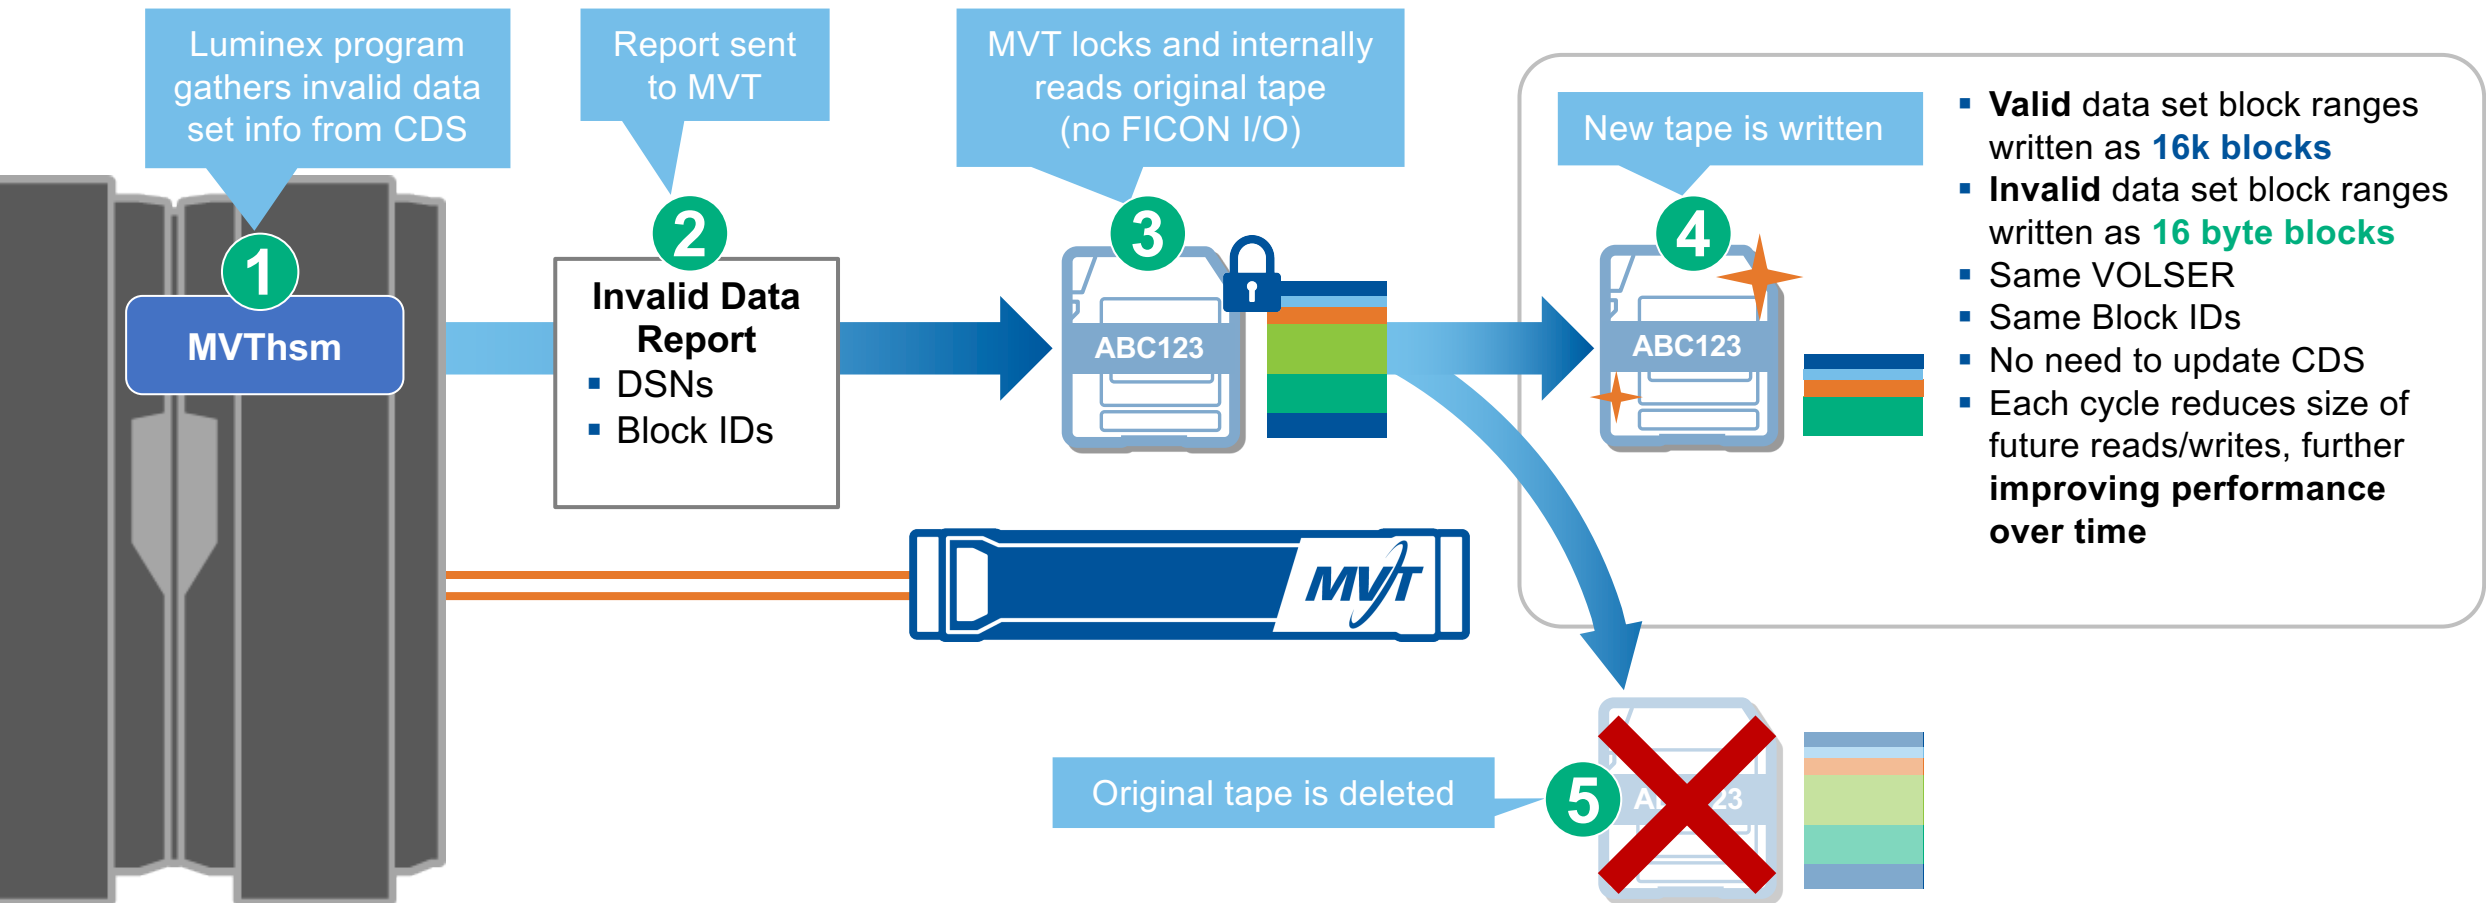


MVThsm: Off-Host HSM Recycle

- Optimize HSM ML2 tape capacity off-host
 - No FICON I/O, done entirely within the tape system itself
 - Reduces expired data set space usage by 99.9%
- VOLSER and block IDs remain the same
 - No need to update the HSM Catalog Data Set
 - The **most CPU-intensive aspect** of the HSM Recycle process
- Performance continues to improve
 - Tapes get smaller with each cycle
 - Faster cloning process
 - Smaller tapes replicate faster
- Effectively reduces tape storage capacity requirements
 - Reduced reliance on mainframe resources = more frequent optimization = less total capacity required
 - Existing users can delay storage upgrades
 - New users can move existing tape workloads into a smaller capacity/lower cost tape system



MVThsm: Process

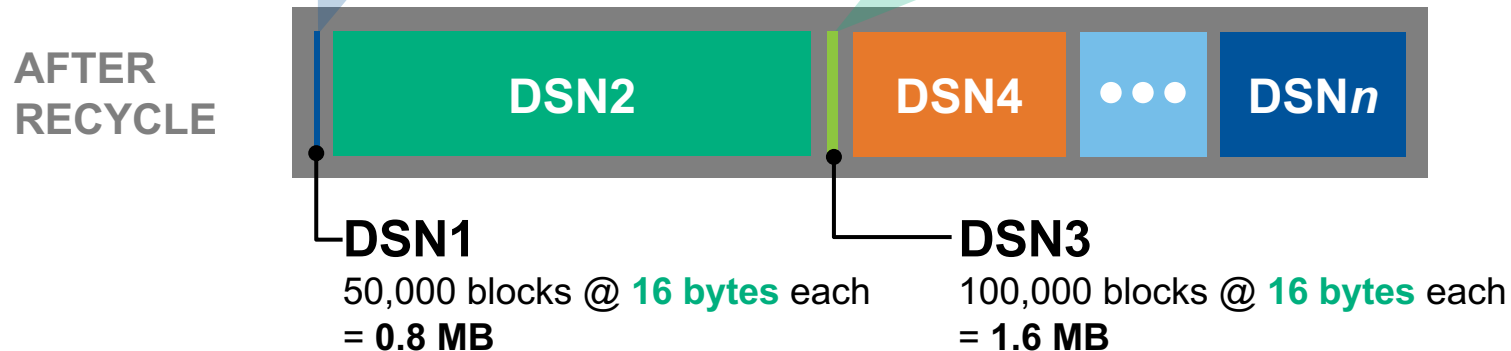


MVThsm: Tape Layout Before and After

VOLSER ABC123: 1,000,000 blocks @ 16 KB each = 16 GB



VOLSER ABC123: 1,000,000 blocks @ 16 KB or 16 byte block sizes = 13.6 GB



850,000 * 16 KB	= 13.6 GB
150,000 * 16 bytes	= 0.0024 GB

Total = 13.6024 GB

- Same VOLSER
- Same Block IDs

Luminex Tape Migration Services

Quick, efficient and non-disruptive tape migration with detailed reporting

Existing Tape Library



Mainframe



Luminex MVT



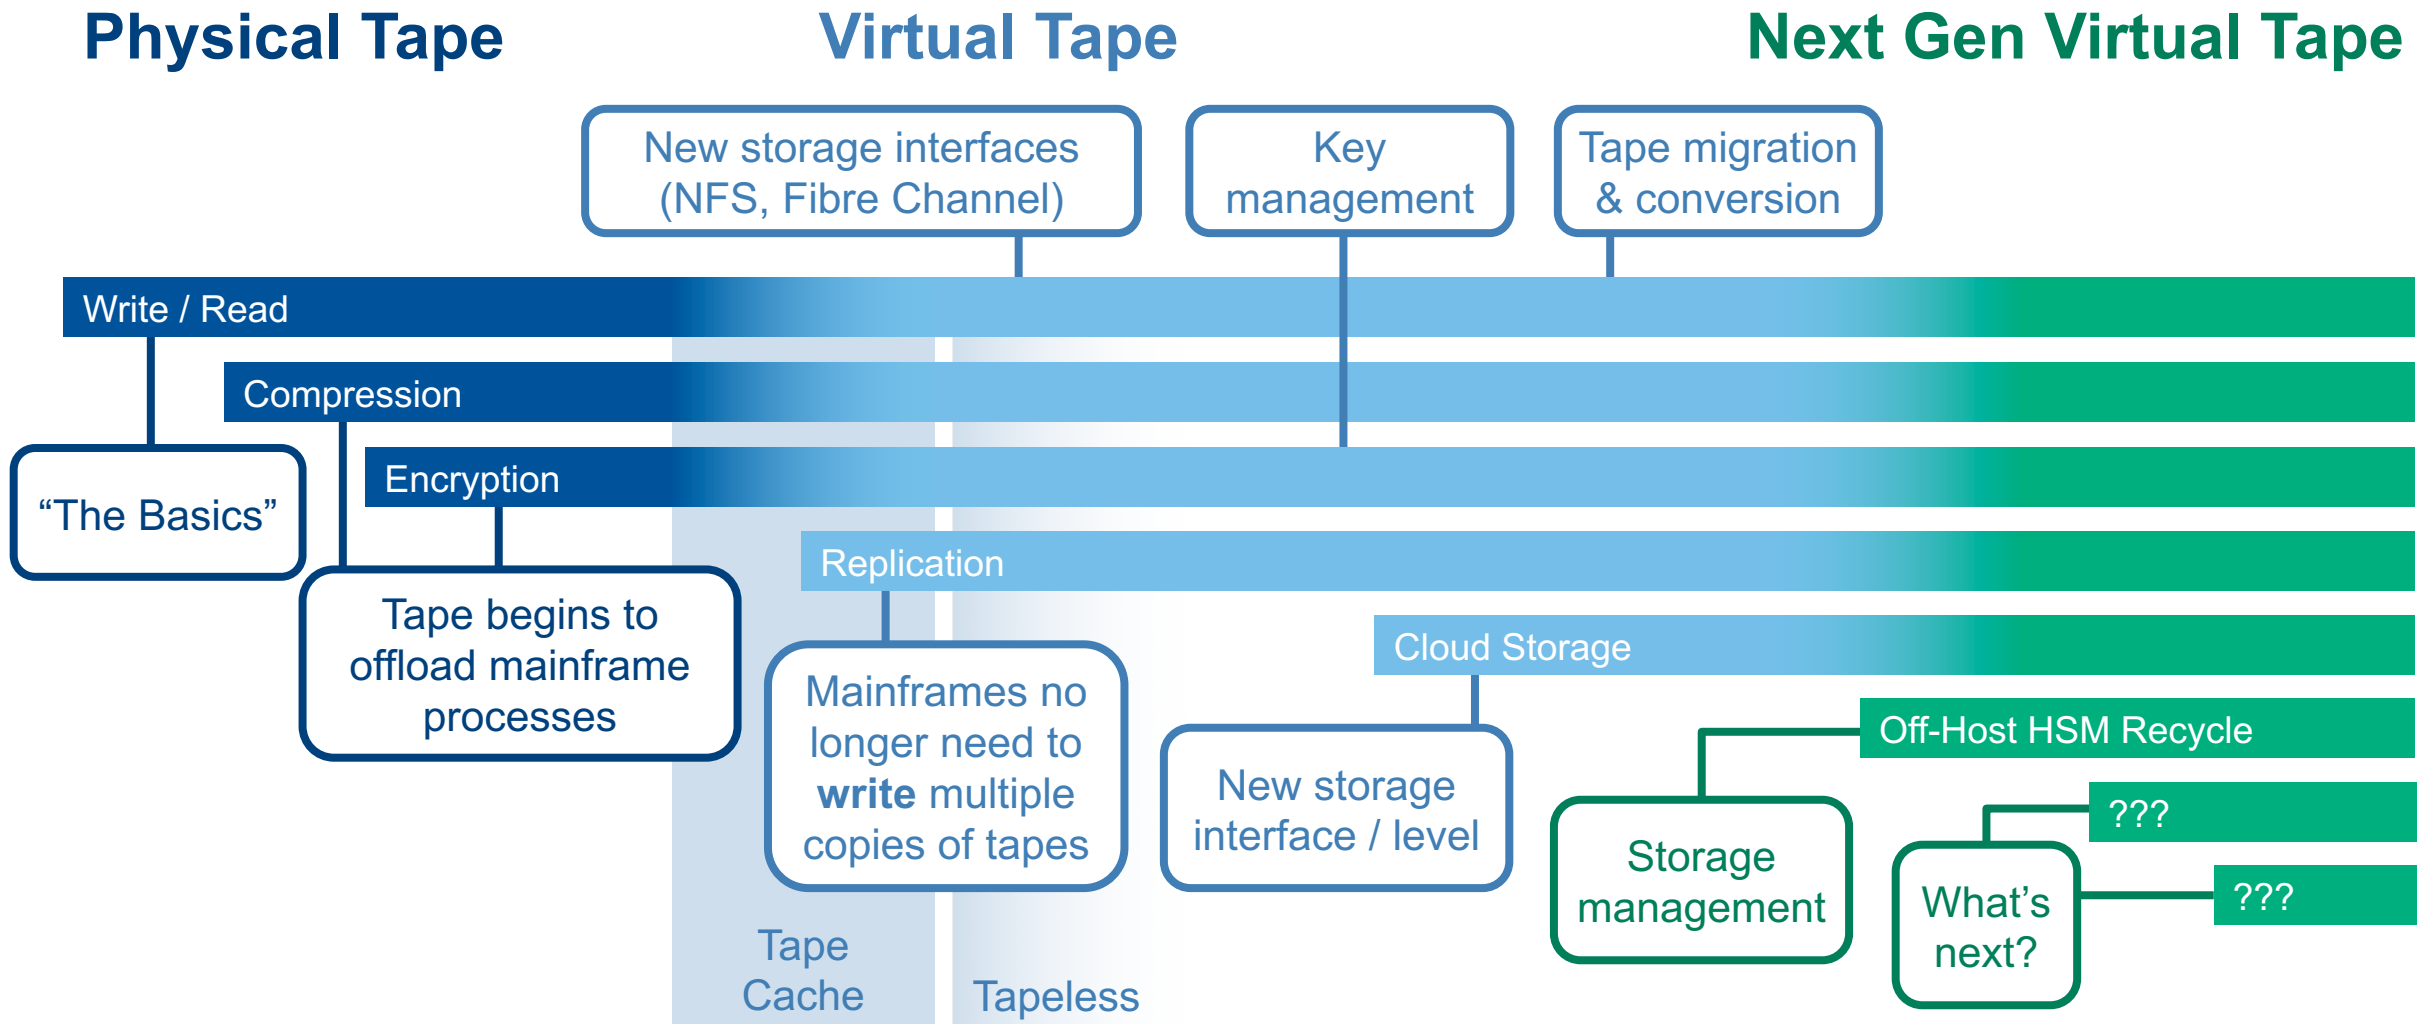
FICON

FICON

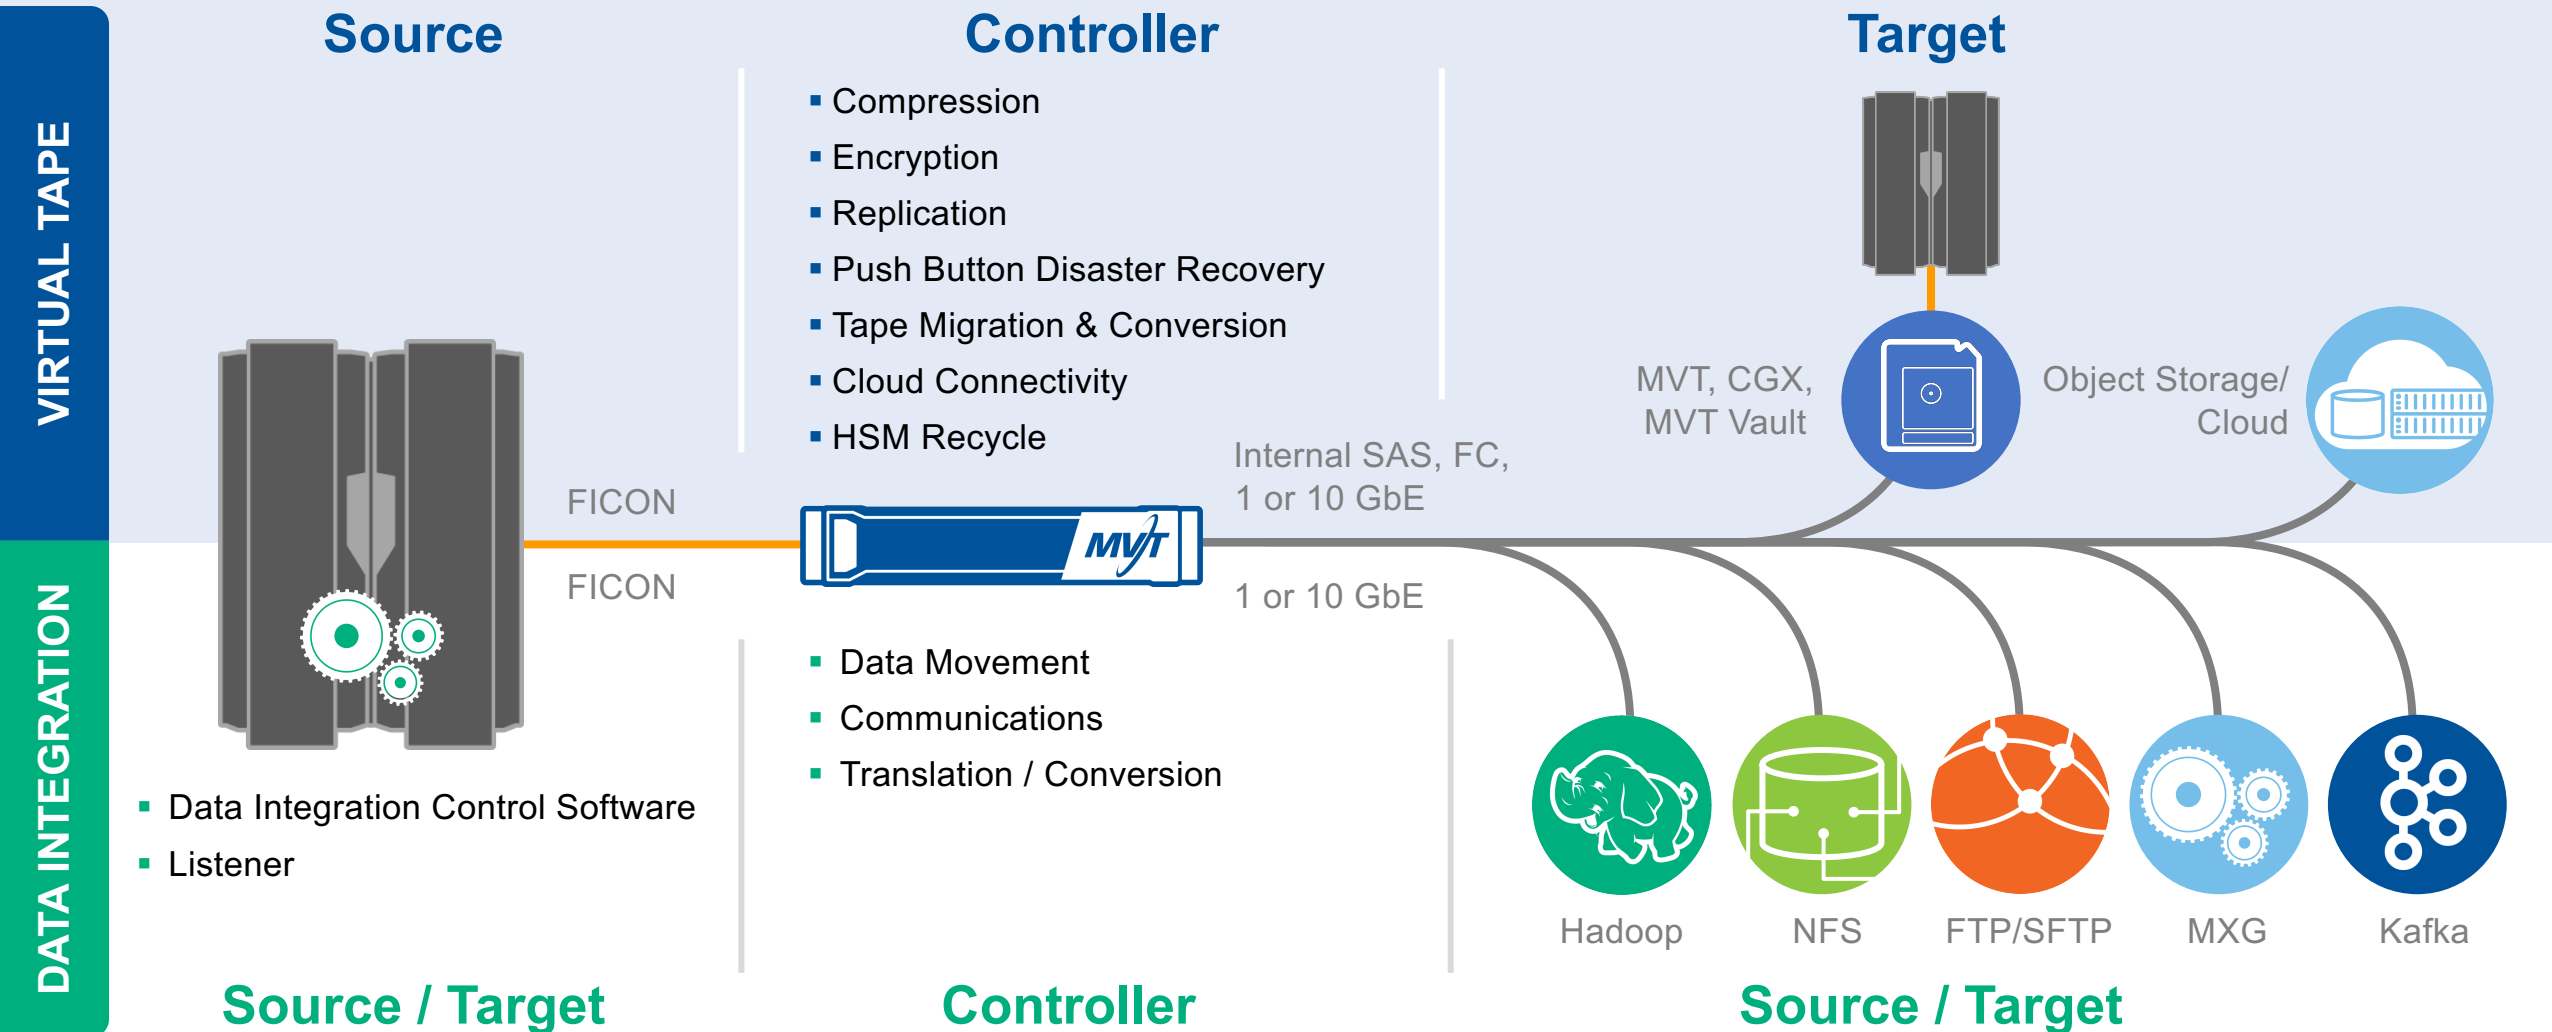
Migration Process

- Assessment
 - Statement of Work
 - Execution
- MIPS-friendly started task
 - Fully compatible with existing MVS or tape management system catalogs
 - Copies one or more tape volsers in parallel
- **Exact copies of original VOLSERs**, including labels, unusual tape mark combinations and any data beyond last tape mark
 - Automatically tracks VOLSERs that have been copied, that have not, and that need to be re-copied
 - Can be ramped up or down

Expect More: Mainframe Tape Continues to Take On New Workloads



From Virtual Tape... *to Virtually Anywhere*



MDI: FICON-based Data Movement and Communications

FICON was designed *specifically for the mainframe*

Attributes

- Fast
- Efficient
- Secure

Uses



DASD



Tape



File Transfer

SecureTransfer: Using FICON as a Secure Data Path

“If you replace mainframe FTP with a channel/**FICON** based solution, you can mitigate FTP security issues a great deal, if not remove them completely. This is the real benefit of a solution such as MDI SecureTransfer.”




Mark Wilson
Technical Director
RSM Partners
www.rsmpartners.com



Top File Transfer Challenge: SPEED

- Faster Data Transfer Rate
 - The FICON connection between the mainframe and the MDI platform copies large data files off the host faster and more securely than TCP/IP based protocols



- Large files are no longer an issue


FICON vs. TCP/IP: Benchmark Testing

30 MB File

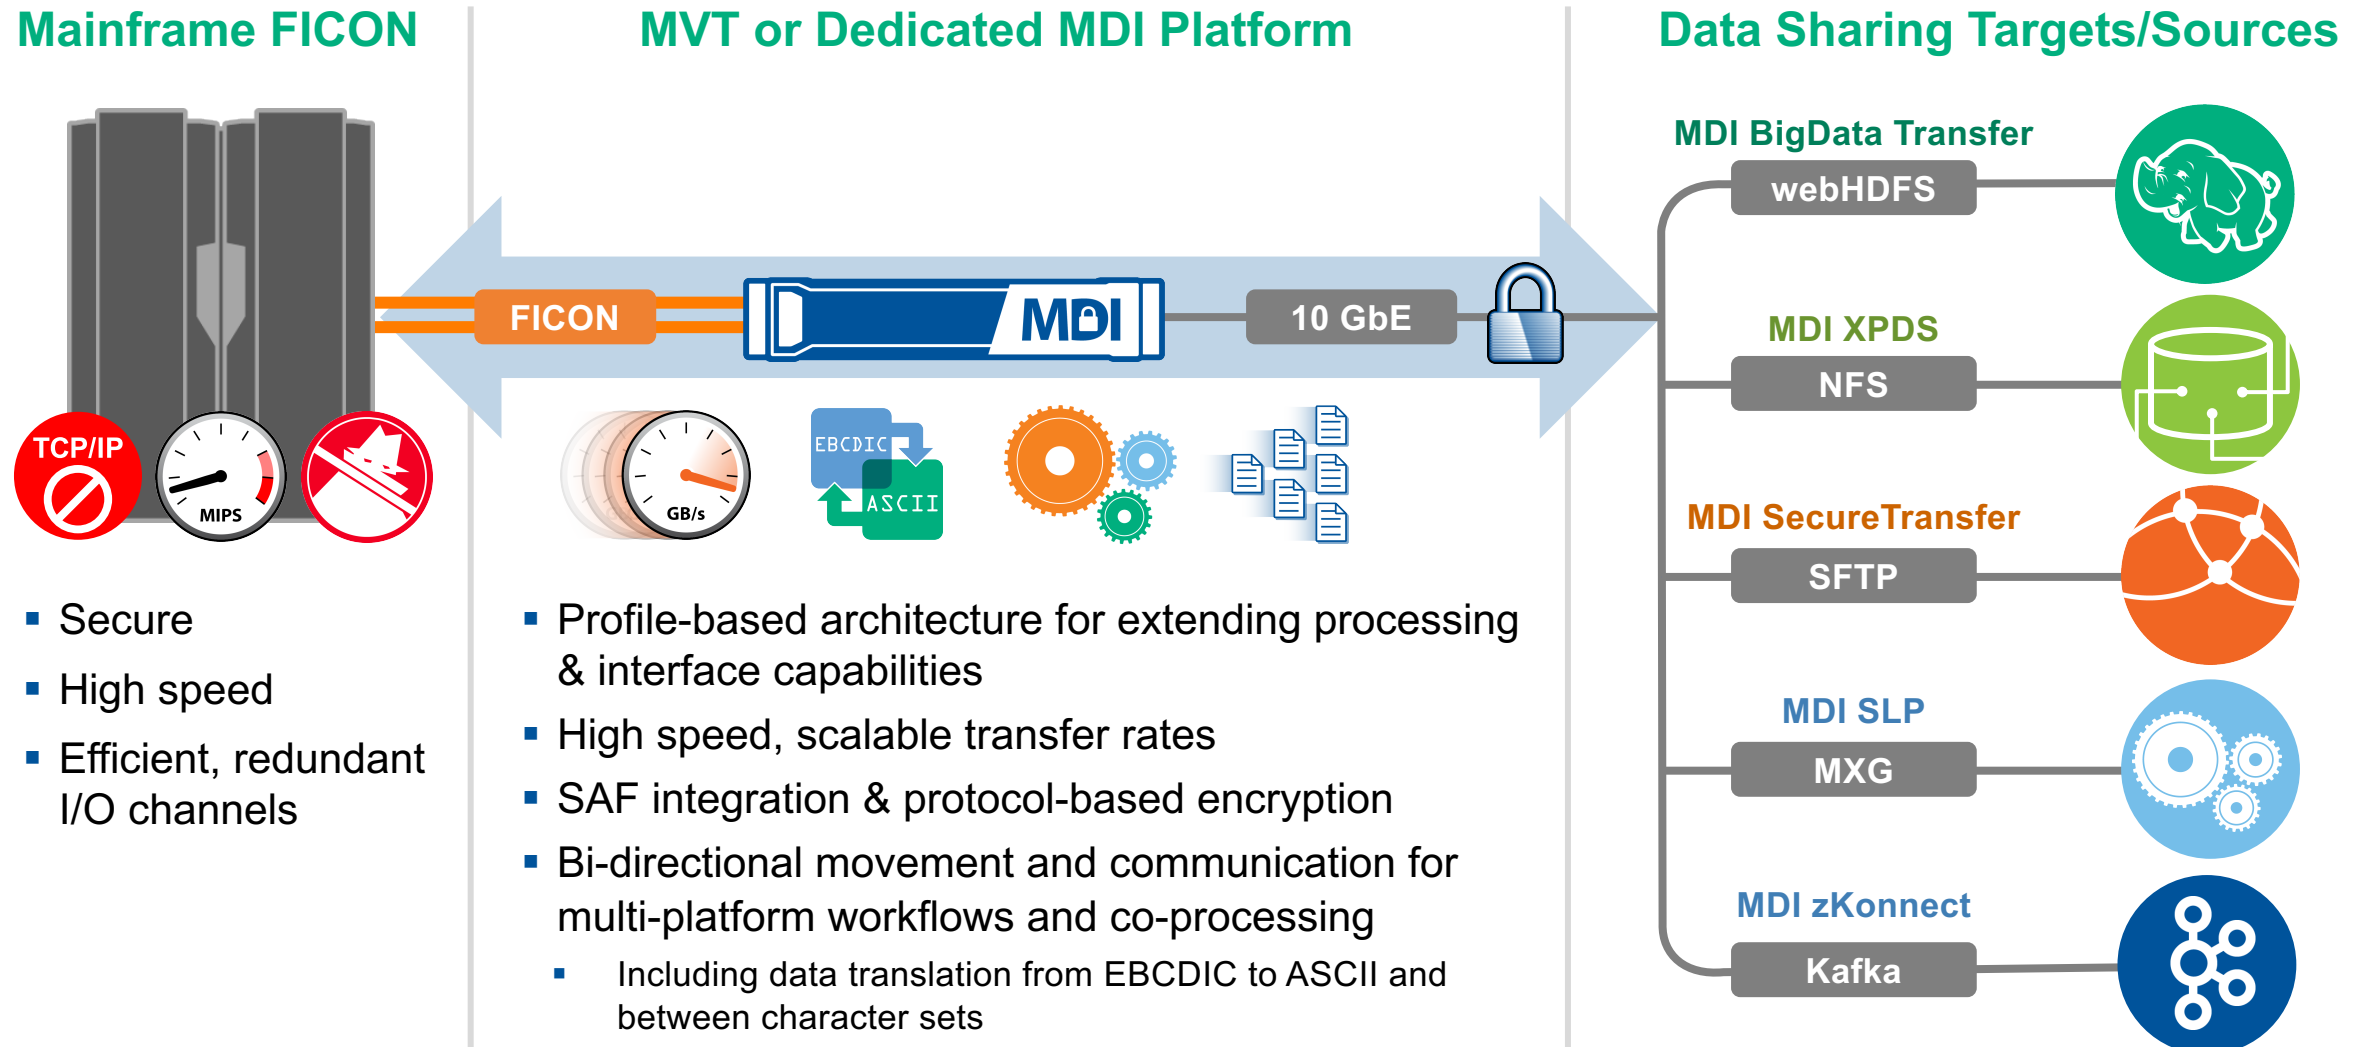
Method	Job	Program	Elapsed	Service Units	CPU
FTP from DSN	BNCHMRK1	FTP	0:00:15.32	2403	1280
(Clear Text)			0:00:15.32	2403	1280
FTP from USS	BNCHMRK2	FTP	0:00:13.96	3060	1409
(Clear Text)			0:00:13.96	3060	1409
SFTP	BNCHMRK3	login	0:00:00.10	150	135
(Encrypted)	BNCHMRK3	tty	0:00:00.02	140	119
	BNCHMRK3	sftp	0:00:00.14	340	317
	BNCHMRK3	ssh	0:00:06.27	68463	34493
	BNCHMRK3	sftp	0:00:08.41	6106	2363
	BNCHMRK3	SH	0:00:08.47	213	163
	BNCHMRK3	BPXBATCH	0:00:08.77	129	107
			0:00:32.18	75541	37697
MDI/IEBGENER	BNCHMRK4	IEBGENER	0:00:03.24	2010	1407
	BNCHMRK4	LUMXPROC	0:00:09.34	156	134
			0:00:12.58	2166	1541
MDI/ICEGENER	BNCHMRK5	ICEGENER	0:00:00.79	667	550
	BNCHMRK5	LUMXPROC	0:00:09.19	151	131
			0:00:09.98	818	681

Benchmarks performed on z13 Model 2965-N10 using SMF Type 30 records

MDI System Resources Savings

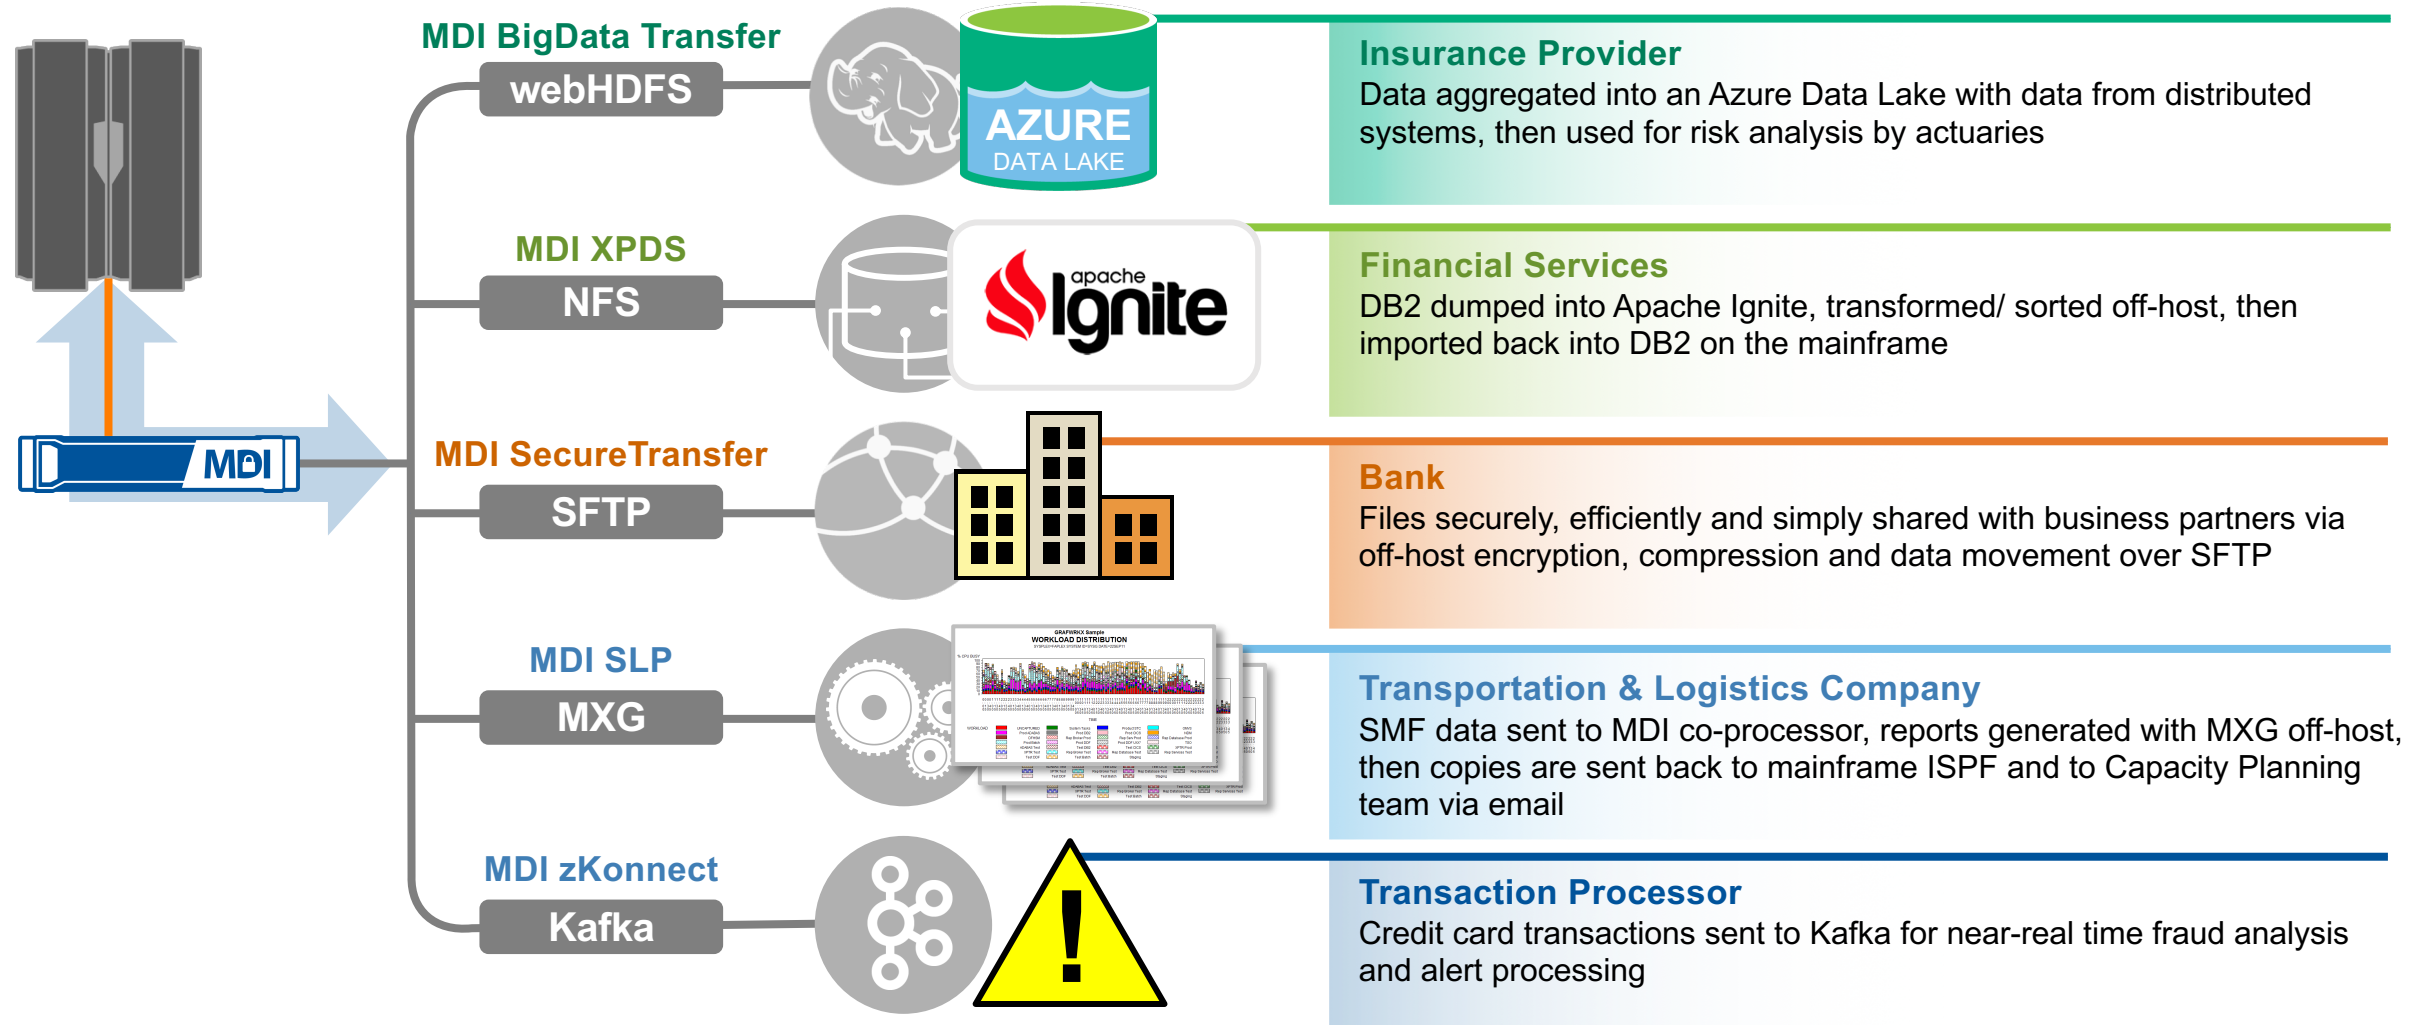


MDI is a Data Transfer & Co-Processing Platform

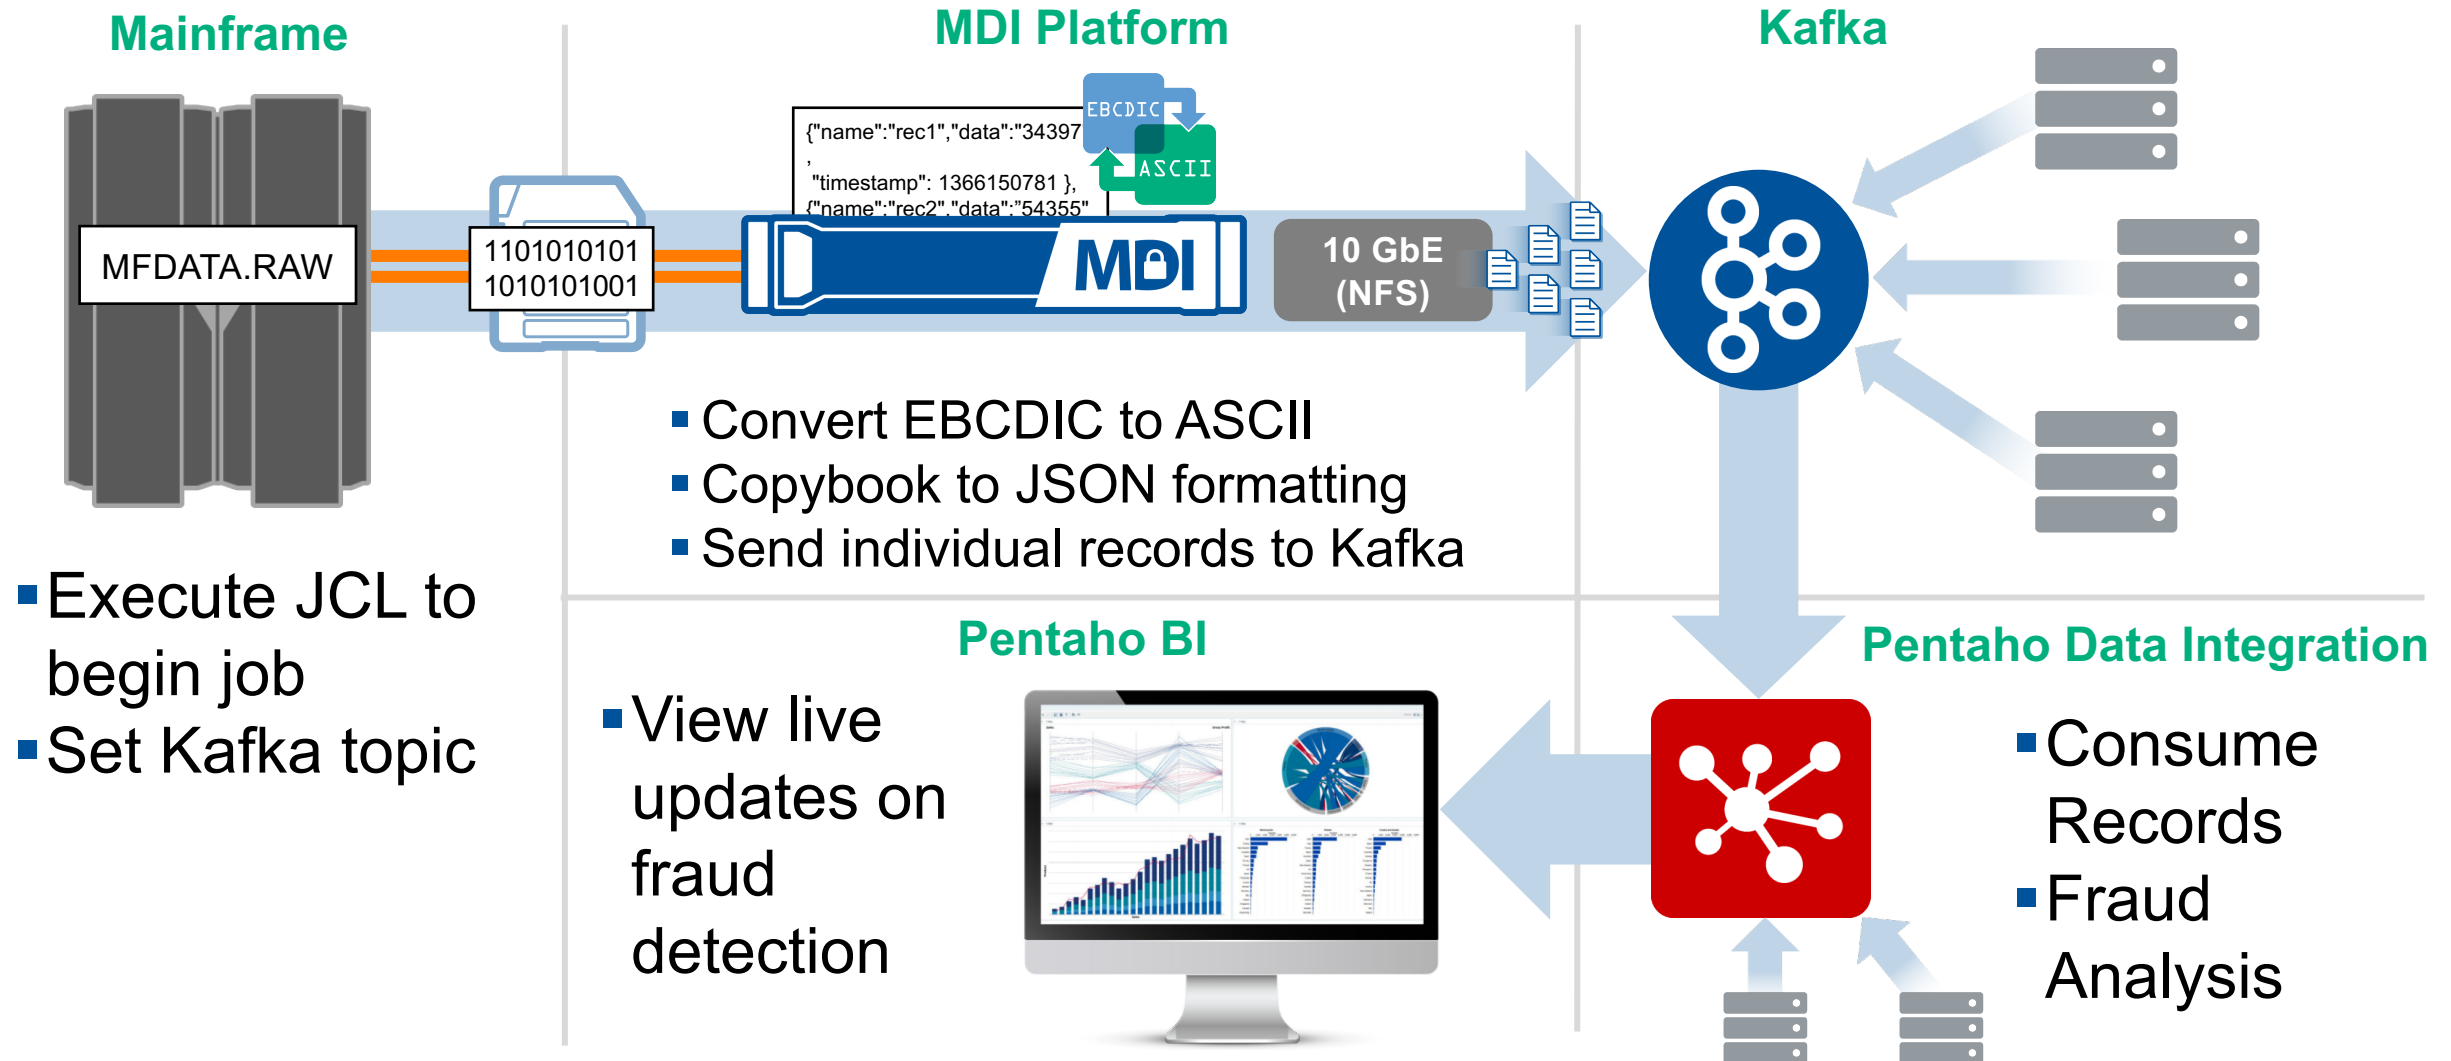


How MDI is Used

SAMPLE USE CASES



MDI and Pentaho: Streaming with Kafka



What's Next

- Add Near-real-time processing in addition to batch with zKconnect Mainframe Kafka Producer
 - LUMXPROC - Batch output written to Kafka (tape volser vmethod)
 - XWRITER - JES2 spool files (output data) written to Kafka in near-real time
 - Log Receiver - System and Application log streams to Kafka in near-real time
 - SMF Receiver - SMF log streams sent to Kafka in near-real time
- Customer-driven integrations for off-host processing such as MDI with Hitachi Pentaho Analytics Platform
- Expand and Secure Integrated Processing to complex mainframe Actions that are not REST API enabled

MDI: A Better Alternative for Mainframe Data Movement



Secure

- More secure than TCP/IP on the mainframe
- Reduce/eliminate open ports on the mainframe
- SAF integration



Fast

- Move data where, when and as often as you want it
- Unmatched transfer rates, scales to the largest data centers
- No bottlenecks or need to “time shift” workloads



Better Business Analytics

- No more silos
- On-demand access to data
- Include “Big Value Data” in enterprise-wide analytics, including Azure



Efficient

- Reduce CPU overhead for mainframe TCP/IP
- Reduce CPU overhead for encryption/translation



Cost-Effective

- Avoid expensive mainframe upgrades by offloading workloads
- Reduce software licensing costs
- Licensing not based on MIPS/MSUs

Luminex MVT: Expect More from Mainframe Virtual Tape

PROTECT

- **STM**: Continuous availability for tape
- **Replication**: One-to-many, cascading, cloud
- **Push Button DR** with non-disruptive testing
- **CloudTAPE**: Archiving and virtual air gap
- **CGSafe**: Encryption & key management
- **MVT Vault**: Standalone replication target
- **SecureTransfer Express**: Easy and secure alternative to mainframe FTP (Port 21)

MANAGE

- **RepMon**: Replication monitoring & reporting, including cloud versioning, tiering status & metadata
- **MVThsm**: Reclaim ML2 tape capacity without mainframe MSUs
- **Multitenancy**: Manage multiple, separate tape workloads at a global or tenant level

LEVERAGE

- **MDI Ready**: Off-host data movement and co-processing platform for a “Connected Mainframe”
- **SecureTransfer Express**: Managed file transfer between mainframes and distributed systems
- **DataStream Intelligence**: Metadata in the cloud

- 

--

1 2 3 4 5 6 7 8 9

1 2 3 4 5 6 7 8 9

1 2 3 4 5 6 7 8 9