

# SMARTDATA FABRIC<sup>®</sup> COMPARISON WITH CONVENTIONAL DATA VIRTUALIZATION AND DATA LAKE QUERY ENGINES REVISION 1.1

## Introduction

There is a number of conventional data virtualization products and high-performance query engines (many federated) on the market, many based on open-source software, e.g., Presto, and, primarily, cloud-based. Some can be installed and run on non-cloud platforms, e.g., on-premise and data centers, and some can run across multiple clouds.

WhamTech develops data virtualization, federation and integration software and has also embraced open-source software for several functions including data discovery, data processing, data security and indexing and query processing. **The main differentiator between WhamTech and all other data virtualization software and data lake query engines is distributed indexing and query processing with a security-centric management system for data, master data and virtual graph data.** WhamTech does not want to develop components that are readily available and accepted in the market, and have a large and active developer base, therefore, WhamTech provides a distributed data virtualization management system for data, master data and virtual graph data that complements many of the conventional data virtualization products and open-source federated high-performance query engines, i.e., a data fabric/data mesh.

### This article compares:

- a. Conventional data virtualization with federation, aka conventional federated data access,
- b. Data lake + query engine, which can be used to populate a data warehouse and/or data marts, and
- c. WhamTech SmartData Fabric<sup>®</sup> that addresses data source and data issues through distributed data management, master data management and virtual graph management, using data virtualization with federation for virtual data integration.

## Five Major Drivers for WhamTech

WhamTech SmartData Fabric<sup>®</sup> is a scalable data fabric/data mesh that is agnostic to query engine, storage, file formats and platforms, e.g., on-premise, data centers, SaaS/third-parties and clouds.

There are five major drivers for WhamTech market acceptance:



1. There is a large gap in the data and analytics market that does not address data issues upfront through data management and master data management BEFORE data is stored/landed in a data lake.

Conventional data virtualization can address data issues AFTER data is stored and landed in a data lake. Query engine vendors are focused on scalability and high-performance, typically, for analytics, but do not address data issues. Query engines tend to assume that data issues will be addressed before landing in a data lake, or in a post-query process, such as ETL to a data warehouse or ELT in a data lake, which causes considerable work for data engineers and analysts to sift through and prepare data.

Data fabrics, and a recent organizational-but-non-technical evolution called data meshes, are tasked with making data available in a highly consumable form through multiple APIs, views and business objects. Some data virtualization vendors, such as WhamTech, describe their software as data fabrics.

WhamTech is focused on delivering high quality, standardized results on-demand, to analysts and normal business users, through what WhamTech calls an Active Data Catalog. WhamTech addresses data (and data source) issues, as well as support for high-performance query processing through indexed-based federated adapters, in the past, using its own index and query engine. Recent advances in open-source query engines and related standards, such as Apache Iceberg, now offer an alternative to indexing associated with database engines such as PostgreSQL, which WhamTech also uses.

2. Not all customers want to, or can, copy all data to yet another data lake or large data warehouse.

An example is a large bank that has a number of data lakes and data warehouses, which are, in effect, data silos, and they want to:

- a. Leave data where it is, i.e., in data lakes and data warehouses, as well as, eventually, operational/transactional systems, therefore, no more or limited data lakes or data warehouses.
- b. Have an **enterprise data catalog** for analysts and business users.
- c. Access data where it is, but subject to **centralized data governance**, and permissions-based access control and data security.
- d. View data and business objects that are standardized, understood and integrated, to support customer 360-degree views and business user-level self-service reporting, BI and analytics.
- e. Implement a solution that is **scalable and high-performance**, with the goal of running most of the compute in the cloud data remains where it resides, aka Hybrid Cloud 2.0, but, in addition, allow for Hybrid Cloud 1.0, where compute may be distributed data remains where it resides.



Other use cases include companies with on-soil data retention regulations, healthcare network providers' ownership of patient data, data security concerns and liability, GDPR, CCPA and other compliance regulations, and SaaS/third-party data sources and data.

### 3. There are no middle-tiers for independent edge data and query processing across multiple disparate data sources

(conventional DV/query engines are only two-tier) to:

- a. Address data source and data issues.
- b. Off-data source (almost no load) uniform and multilateral high-performance SQL query processing.
- c. Create and maintain near real-time derived values and pre-aggregated, pre-calculated and pre-joined views.
- d. Provide lower-level federation for department/division/LOB reporting, BI and analytics.
- e. Support DATA MESH data-as-a-product as part of a data fabric.

The benefits would be more data intelligence/metadata, query acceleration, and reduced network load and latency in transferring large amounts of raw data to a data lake and/or query engine.

Current conventional data virtualization with and without federation, and federated query engines, both tend to be two-tiered – data source systems and query engine. This two-tier architecture includes source data that could be copied and landed to a centralized or distributed data lake, which is just a change in location of the data source.

- 4. Advanced enterprise-wide access control, and data governance and security that operates on ALL data sources and data, regardless of data source support.
- 5. Little or no support for event processing and interoperability, and other use cases that are not just analytics.
  - Real-time data access and monitoring.
  - Read AND WRITE data.



#### "Data Fabrics for Frictionless Data Access"

According to Rick F. van der Lans, Independent Business Intelligence Analyst, R20/Consultancy, April 2021, IT systems must provide the following thirteen capabilities:

- 1. Data preparation, such as transformations, aggregations, filters and joins
- 2. Adaptable logic
- 3. High-performance
- 4. Data access by many data consumption forms
- 5. Access to all the data sources
- 6. Processing of all types of data
- 7. Data security and privacy
- 8. Real-time data access
- 9. Read and write data access
- 10. Data minimization
- 11. Event processing
- 12. Technical and business metadata management
- 13. Master and reference data management

SmartData Fabric<sup>®</sup> provides the above scalable data fabric capabilities, data management and master data management that is agnostic to high-performance query engines, data sources, types, formats, locations or platforms, and run from a cloud, on-premise or hybrid cloud.

The following table delves into details on the above thirteen capabilities by defining them and comparing conventional data virtualization (DV), data lake query engines and WhamTech SmartData Fabric<sup>®</sup>:



Capability				Data lake query engine	
	-			arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
1	Data preparation,	Data formats, quality and	Almost completely dependent	Most DV vendors recognize	Deals with data issues BEFORE any
	such as	standards vary	on data sources for data	this problem and recommend	queries on data sources are made. Uses
	transformations,	considerably, sometimes	formats, quality and standards,	that DV find and read data	near real-time indexes to (1) profile,
	aggregations,	in same systems, but	and in many cases, cannot pre-	from sources to populate a	identify, classify and match data, and
	filters and joins	almost always across	process data such as	data lake. Once landed, data	develop and test data transforms to
		multiple systems, e.g.,	aggregations, filters and joins.	formats, quality and standards	cleanse, transform and standardize data
		mainframes, Big Data and	Different names of columns/fields	can be addressed by creating	used to build and maintain indexes, and
		social media. Many times,	are not usually a problem, as can	materialized views on top of	on results data, (2) support Forrester Zero
		raw data needs to be	be mapped to a standard, and	data, or subsequent ELT	Trust Data Security Framework, (3) pre-
		cleansed, transformed	columns/fields can be separate	within the data lake or ETL to	process data, e.g., derived values and text
		and standardized,	terms in a query, but more difficult	a data warehouse.	analytics, (4) link data within and across
		enhanced and	if needed to be split up, e.g., last		data sources, (5) provide multiple virtual
		aggregated, filtered and	name as part of single column first		standard views, (6) allow full data
		Joined within and across	and last name. However, the		traceability, (7) execute high performance
		data sources, for it to be	success of external queries is		distributed parallel queries, (8) monitor
		a readily consumable	completely dependent on source		data sources, (9) enable virtual graph
		product.	data formats (type and size) and		database link analysis and visualization,
			data quality. This yields potentially		and (10) invert and generate results when
			incomplete and incorrect results		the data source is N/A.
			data, which can also be		
2	A devete la la sia	Different data anymera	Uncompliant, e.g., HIPAA and PCI.		Mant data in an antianal (transportional
2	Adaptable logic	and different types of	Similar to capability #1, external	Again, similar arguments as in	wost data in operational/transactional
		data require different	dependent on data sources and	capability #1, whereby	it is non-human generated however for
		approaches and these	data formate quality and	lake provides the ability to	the data to be useful and immediately
		can change over time	standards. Once data is initially	apply different logic to the	consumable, it may need to be
		can change over time.	read it can be used to populate	data as needs change	aggregated calculated or joined with
			various views and many data	data as needs change.	other data within the data source and/or
			issues can be addressed and		across multiple data sources. There may
			updated over time. However, the		also be a requirement for derived value
			initial guery and read on the data		indexes and indexed views Different
			source is needed, and that is		types of indexes and indexed views allow
			where the problems occur.		great flexibility and adaptability over time
					as both data sources and business
					requirements can change over time.



Сара	bility			Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
3	High-performance	This generally refers to fast query processing and	Similar to capabilities #1 and #2, query processing and results	Similar, to capabilities #1 and #2, data lake engines	Use of independent middleware-based independent external indexes and
		results retrieval.	retrieval performance are	recommend copying all data	indexed views, allow for high
			dependent on data sources. In	out of source systems to	performance and more intelligent
			some cases, source systems are	overcome any query	external queries. Intelligence comes
			high performance and can sustain	processing and query load	from pre-aggregated, calculated and
			a demanding external query,	issues.	joined views, and derived value
			however, in general,		indexes, avoiding large raw data source
			operational/transactional	Many data lake engines do	query processing load and times. There
			systems are not designed to	not address data latency, but	may be cases where data source reads are
			optimally process reporting, BI	some do and introduce	necessary for intermediate results or final
			and analytics-type queries. Even	changed data capture either	results data, but these are low-level,
			performant systems can	directly from the sources or	usually sequential reads that have a low
			experience a significant external	though frequent query	impact on source systems. Alternatives
			query load impacting	polling.	include maintaining data in indexes.
			operations/transactions.		
				Data lake engines tend to use	
				II "horsepower" to address	
				performance vs. intelligence,	
				which comes at a very high	
				cost.	



Сара	bility			Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
4	Data access by	Multiple APIs and means	Conventional DV can	Similar to capabilities #1, #2	WhamTech index and query engine
	many data	to access the same data,	accommodate data access by	and #3, data lake engines	and associated data management and
	consumption	depending on the	many data consumption forms.	suggest that is it is easier to	MDM, is the same, regardless of the
	forms	application, and unlike	But, uniformity of SQL query	build and maintain many	different types of data sources or data
		operational/transactional	processing, views, business	uniform data consumption	in these sources either in their original
		systems, there may be	objects and APIs across multiple	forms in a centralized or at	locations or a data lake. This allows for
		multiple different	disparate data sources, may be	least centrally managed	centralized management, and "cookie-
		applications and users.	difficult to achieve and manage	distributed data store.	cutter" views and data consumption
			because of the different types of	However, the challenges of	forms, approaching the benefits of a
			data sources and data in these	dealing with different types of	centralized or centrally managed data
			sources.	data sources and data in	warehouse - virtual.
				these sources, remain. EIL in	
				a data warehouse, ELI with	
				management system will still	
E	Access to all the	There is a broad	As conventional DV is almost	De needed.	Whom Tash propresses data from a
5	data sources	spectrum of data sources	As conventional DV is almost	dool with all types of data	broad range of data sources as it builds
	uala sources	available some of which	sources and data in sources for	sources in controlized data	and maintains indexes and indexed
		may not be open to	external queries, they typically	lake This is an advantage in	views, which then are subject to the
		external SOL type	provide connectors and query	the sense that data in a data	same uniform SOL queries as all
		queries There either	conversion for non-SOL data	lake regardless of source	adapters to other data sources
		needs to be a query	sources and in some cases they	may become queryable	WhamTech has had a lot of experience
		conversion or a	may build and maintain cache or	depending on the data lake	dealing with difficult data sources e.g.
		quervable abstraction to	SOI -compatible external views –	platform. Some data sources	IBM Mainframe Cobol-generated VSAM
		accommodate all types of	virtual or materialized. Some data	may require schema-level	files, and social media feeds. The indexes
		data sources.	sources are just not capable of	transformation to a data lake	and indexed views provide a SOI -
			any kind of eternal query.	platform.	compatible abstraction of source data.



Capability				Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
<b>No.</b> 6	<b>Description</b> Processing of all types of data	Definition Data types, formats, quality and standards vary considerably, sometimes, within the same data sources, but certainly across multiple data sources. A Data Fabric has to be able to accommodate these considerable variations.	Virtualization/Federation ("DV") Because conventional DV is almost completely dependent on data sources and data in sources for initial query processing, it has tended to only accommodate queryable structured data sources. Some conventional DV vendors may deal with unstructured data, but not until it has been read from data sources as part of a results set with queries on structured data.	centralized data store Once data from sources is landed in a data lake, there would still be a need to process data to address data issues, create derived values and views, and run text analytics, e.g., entity extraction, meaning and sentiment. One advantage of a data lake over conventional DV would be that queries could be made directly on unstructured data. More difficult data types would still require data processing to	WhamTech DVAs mentioned in previous capabilities,WhamTech cleanses, transforms andstandardizes data as it builds andmaintains indexes and indexed views.That data can include unstructured datathat requires preprocessing, text analytics,etc., and perhaps other difficult data typesand formats.WhamTech provides a uniform datasource access and data, regardless ofdata source type or data types.
				transform and standardize data, in addition to	
				addressing data quality and MDM.	



Capability				Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
7	Data security and	Provide enterprise-wide	Conventional DV has such	Again, as per previous	WhamTech addresses data issues
	privacy	data governance, access	capabilities and as such, can be a	capabilities, data lake engines	upfront, even automatically identifying
		control and data security.	gatekeeper to data. But the	make the argument that	and mapping data to standard data
		Data Fabrics should be	problem comes in being able to	centralized data governance,	views, subject to data engineer review.
		able to take advantage of	define what data is and	access control and data	By addressing data issues upfront, this
		centrally managed	standardize it to allow these	security would be easier,	allows data governance, access control
		LDAP/Active Directory	capabilities to be automatically	which would be true.	and data security to be automatically
		Domain Controllers and	applied. There is a lot of manual	However, there are still the	applied.
		associated user	work involved as data engineers	issues conventional DV	
		credentials and ID	attempt to understand the data	struggles with, namely the	WhamTech also supports dynamic data
		Authentication, SSO,	for data governance and,	manual work involved by data	masking, tokenization and encryption,
		RBAC, ABAC/RLS and	perhaps, map it to a standard	engineers and the potential	depending on user credentials, as per the
		CLS. This may also	view or policy document. Then,	compliance issues.	definition.
		involve working with	due to data type or quality,		
		Virtual Directories across	there may be questions about		
		multiple domains, and	incomplete or incorrect results		
		using distributed options	data, which would be		
		such as SAML.	compliance issues, e.g.,		
			CCPA/CCPR, GDPR, HIPAA and		
			PCI.		



Capability				Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
8	Real-time data	Access data as it is in	<b>Conventional DV does this very</b>	Data lake engines tend to	WhamTech DV is similar to
	access	sources – no or minimum	well, accessing data as it is, in	downplay this capability as	conventional DV in that it enables real-
		latency.	sources, because it is almost	they say their primary	time data access. Typically, indexes and
			completely dependent on data	purpose is to provision data	indexed views are updated in near real-
			sources for query processing.	for exploration and analytics –	time, using one of at least twelve,
			Dependency on data sources and	not operations. They also	changed data capture (CDC) methods
			data in sources have other issues	propose that analysts need	from data sources.
			discussed in other capabilities, but	snapshots of data to work	
			it does enable real-time data	with, so, real-time data is not	Most operation/transaction systems do
			access.	desirable. A few data lake	not change the bulk of their data, as
				engines have implemented	operation/transaction data is usually
				real-time updates to data	appended, not changed or updated.
				lakes – see comments under	Some changes/updates will occur in entity
				capability #3. The real-time	data – customer, and organization info,
				update argument is that the	etc. Reference data can also change over
				data lake is an active data	time, e.g., prices, exchange rates, etc.
				catalog and can be used by	These entity and reference data changes
				data engineers and analysts	can be tracked and retained in master
				to create data extracts that	data and reference data management.
				can be used and archived	
				independently to run	
				analytics.	



Capability				Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
9	Read and write	Provision of APIs and	Most conventional DV is	Data lake engines are typically	WhamTech DV supports write data
	data access	other forms that allow	intended for reporting, BI and	not intended for	access. The same constraints as
		both read and write data	analytics, which tends to be	operations/transactions or	described in the definition apply to
		access. Reading data is	read only and do not attempt to	interoperability, but that is	WhamTech.
		considerably less difficult	write data access. There may be a	changing as more enterprise	
		than writing data, as	few exceptions.	systems migrate to the cloud.	The major benefit of WhamTech DV is
		writing data requires		As in capabilities #1, #2 and	that the standard data views apply to all
		transaction management	Data fabrics that include write	#3, data lake engines argue	data sources, so, updates, such as
		through stored	data access extend the role and	that copying all data to a	customer info – email, phone numbers,
		procedures, or an app (of	capabilities of conventional DV to	centralized data lake is the	addresses, etc., can be universally and
		some form) to check and	a much broader range of use	best option for read data	uniformly understood and applied. Data
		potentially update	cases, including operations.	access, but while there are	may have to be reverse transformed back
		associated data in other		ELT (extract, load and	to the original data type and format.
		tables, for example. The		transform) capabilities	
		degree of desired		mentioned. In reality, most	
		interoperability will		companies will ETL to a data	
		determine how involved		warehouse or data marts of	
		the write data access		some form.	
		process will be. The good			
		news is that systems are		Data lake engines are not set	
		increasingly making write		up to support write data	
		APIs available with built-		access or interoperability on	
		in stored procedures and		original data sources.	
		transaction management.			



Сара	bility			Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
10	Data minimization	Avoid making multiple copies of data where possible. The ratio between unique and replicated data will be 1:10 by 2024. That ratio was 1:9 in 2020. (Source: IDC). This also impacts data security, compliance, data latency, cost and knowing what data is correct or not.	The major benefit of DV, in general, is that data remains in sources, thus either eliminating or minimizing data replication, and tending towards real-time data access with no latency.	Data lake engines claim that because of data source and data issues, it is better to accept that one good copy of data is made that can be worked with vs. multiple copies. In reality, even modern data architectures make multiple copies of data, introducing data silos and latency, and other issues mentioned in the definition.	Similar to conventional DV, WhamTech DV eliminates or minimizes data replication and tends towards real-time data access with no latency.
11	Event processing	Depending on key data value changes that can be monitored, subsequent events can be triggered. These events drive processes and can lead to process automation. In effect, event processing allows for data-driven app development that can bridge data sources and enable sub-processes.	Conventional DV cannot directly monitor key data changes occurring in sources. Instead, sources either have to be polled or queries frequently made, leading to data source load and performance issues, and tending away from real-time events.	Data lake engines rarely mention event processing as it is considered operational instead of analytical (reporting or BI).	WhamTech DV uniquely builds and maintains indexes on key data values, typically, in near real-time, that can monitor data source changes. Many KPIS are actually aggregations, calculation and/or joins of source data, but that data has to be curated, i.e., data issues addressed, to be useful. Stored triggers on these values can lead to event processing. WhamTech leverages REST APIs and standardized BPM language and software, primarily, open source.



Capability				Data lake query engine	
				arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
12	Technical and	"Data about data" is the	Through data discovery and	In many ways, data lake	WhamTech DV addresses both
	business metadata	conventional description	sampling, conventional DV can	engines perform the same as	technical and business metadata
	management	of metadata and business	provide and visually present	conventional DV, leaving	during building and maintaining
		objects can be a part.	technical metadata. However,	business metadata attributes	indexes that provide data profiles for
		Technical metadata refers	many times, it is left up to data	up to data engineers, analysts	data discovery, identification,
		to attributes such as type,	engineers, analysts and end-	and end-users. Only if and	cleansing, transforming and
		format, distribution and	users to know what the business	when lake data is ETL'd to a	standardizing data. Before the first query
		stats. Business metadata	metadata for source data are.	data warehouse or data	is made, the end-user/app is presented
		refers to how the data	There is little to no attempt to	marts, would there be	with business metadata, which is
		matches attributes	cleanse, transform and standardize	attempts made to conform to	managed bottom-up instead of the
		associated with business	data to conform to standard data	business metadata.	conventional top-down approach. This
		standards, context and	views and models, and therefore		lends the WhamTech DV solution to
		relationships, of which	business metadata. Once read and		being an Active Data Catalog that can be
		business objects could be	landed somewhere, source data		used by almost any user in an
		a significant part.	can be cleansed, transformed and		organization.
		Standardized business	standardized.		
		metadata is an essential			
		part of data governance,			
		data security and access			
		control.			



Capa	bility			Data lake query engine	
	-			arguments against DV and	
			Conventional Data	for copying all data to a	
No.	Description	Definition	Virtualization/Federation ("DV")	centralized data store	WhamTech DV
13	Master and	Master data describes the	MDM and RDM are, typically,	MDM and RDM are not	WhamTech DV has long-recognized
	reference data	people, places, and	separate systems that are difficult	usually part of a data lake,	that MDM, in particular, and RDM, are
	management	things that are involved	to integrate into normal	and may be considered	essential to data access, data
		in an organization's	operations and reporting, BI and	separate source systems	integration and a data fabric that can
		business. Reference data	analytics. Some conventional DV	within the data lake, or after	support both operations, and
		are sets of values that are	vendors do provide MDM as	source data is landed and	reporting, BI and analytics. WhamTech
		referred to by systems,	part of their platform, but then	part of ETL to a data	MDM is an integral part of its data
		applications, data stores,	usually there are separate	warehouse, or some other	fabric that allows a seamless and
		processes, and reports, as	processes for building and	form of data management.	automatic integration of master and
		well as by transaction and	maintenance, and not an		reference data with data access. The
		master records. Master	integral part of data access and		indexes for master data are combined
		data management	data management.		with content and link data indexes, to
		(MDM) and reference			provide multiple standardized views of
		data management (RDM)			data, including virtual graph data. Master
		are essential to data			data can also be distributed and
		integration, and			associated directly with adapters to data
		providing any value to			sources and coexist with other indexes –
		operational data for			this avoids centralization, eases privacy
		almost any reporting, Bl			and security concerns, meets regulations
		and analytics. Both MDM			and improves performance. WhamTech
		and RDM should be			DV can automatically use master data to
		integral parts of any data			find any associated source data, enabling
		solution, as they enable			entity-based views such as single
		views of how a business			customer/patient views and many more.
		is performing and allow			
		data to be leveraged to			
		provide insights and			
		value. Without integrated			
		MDM and RDM,			
		businesses will not be as			
		successful as they could			
		be and will find it difficult			
		to compete.			



For more technical information, please contact Gavin Robertson, CTO and SVP, WhamTech, <u>gavin.robertson@whamtech.com</u>, 972-991-5700 x706 (o).

For sales and investment information, please contact Mark Armstrong, President, WhamTech, <u>mark.armstrong@whamtech.com</u>, 972-991-5700 x708 (o).

Copyright © 2022, WhamTech, Inc. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. Product names may be registered trademarks of their respective owners.