Infrastructure of Tomorrow: the Compound Data Center Architecture

Michele Del Sordo, Executive Architect & Executive Consultant
December 2016

Today context: the digitization phenomenon and the consequent data deluge

- Our world is changing at a pace that is difficult to comprehend and be aligned with.
  - Global integration – at an unprecedented scale.
  - Digital technology -- becoming increasingly embedded into every object and system.
  - Smarter networking technologies -- connecting people and systems in entirely new ways.

- These are making the world a smaller place. However, something even more compelling for businesses, communities and individuals is happening -- digital data is flooding our society (the Internet of Things phenomenon).
  - According to one estimate, mankind created 150 exabytes (billion gigabytes) of data in 2005. This year, it will create 1,200 exabytes. Keeping up with this flood of data and storing useful data is difficult enough. However, analyzing the data to spot patterns and extract useful information is even more complex. The data deluge is already starting to transform business, government, science and everyday life.

A Data Deluge example: the Automotive case

- Automobiles have a plethora of sensors, constantly streaming data about the car and its environment.
- The data collected is used to provide a number of convenience and safety features to the operator.
- The data could be highly useful to other cars in the area
  - Icy conditions ahead
  - Stopped/slow traffic
- The data could also be useful at a more macro level:
  - Should the town send out additional road crews to treat the road for ice?
  - Is there a traffic incident that needs police attention?
- Even broader long range uses for the data would provide additional value
  - What are the typical traffic patterns based on days of the week/months of the year?
  - In what weather conditions do we need to plan to send out road crews for snow removal and how often?
  - How does one city compare to another in terms of traffic management, snow removal?
- Making use of and managing this data will take a new type of data center compared to current data center topology.

Data in the Automotive case (1/2)

- Vehicle Data
  This includes the various measured attributes and reported from the vehicle based on


INTERNATIONAL STANDARD | ISO 22837

Vehicle probe data for wide area communications

ISO 22837:2009 VehiclePackage																																		SurroundingPackage										RN Pack		
Vehicle.latitude	Vehicle.longitude	Vehicle.altitude	Vehicle.velocity	Vehicle.acceleration	Vehicle.yawRate	Vehicle.stoppingTime	Vehicle.engineStoppedTime	Vehicle.lateralAcceleration	Vehicle.direction	Vehicle.vehicleType	Vehicle.vehicleUsage	Vehicle.Gforce	Vehicle.suddenSteeringManoeuvre	Engine	FuellingSystem.fuelConsumption	FuellingSystem.averageFuelConsumption	DriveTrainSystem	Chassis	BrakingSystem	Brake.status	Brake.boostAssist	AntiLockBrakeSystem.status	TractionControlSystem.status	VehicleStabilityControl.status	ParkingBrake.status	SafetySystem	InteriorEquipment	Wiper.status	SeatBelt.status	Trunk.status	Door.status	OnboardInformationEquipment	ExteriorLights.status	Environment.rainfallIntensity	Environment.temperature	Environment.lightCondition	Road.longitudinalSlopeScale	SurroundingOBject.distance	SurroundingObject.direction	Obstacle.detected	Obstacle.distance	Obstacle.direction	LaneMark.detected	Path.exceptionVariance	Link	Point

- Driver Data: driver's attributes, preferences, etc..
- Geographical Data: map data
- Area Data: specific area data

Data in the Automotive case (2/2)

Cause code
Traffic congestion
Accident
Roadworks
Narrow lanes
Impassibility
Slippery road
Aquaplaning
Burst pipe
Fire
Hazardous driving cond.
Obstructions on the road
Animals on roadway
People on roadway
Broken down vehicles
Rescue and recovery
Regulatory measure
Visibility reduced
Wrong way driving
Intersection violation
Stationary vehicle
Collision risk warning
Emergency vehicle appro.
:

- Event
  Information about traffic events, e.g. road works, traffic jams, etc.

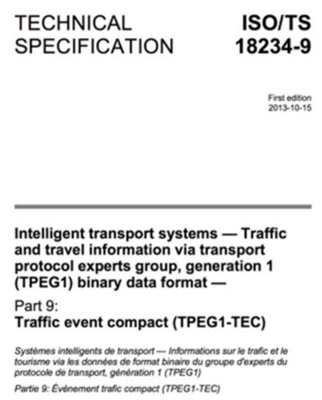
TECHNICAL SPECIFICATION

ISO/TS 18234-9

First edition 2013-10-15

Intelligent transport systems — Traffic and travel information via transport protocol experts group, generation 1 (TPEG1) binary data format —

Part 9:
Traffic event compact (TPEG1-TEC)

Systèmes intelligents de transport — Informations sur le trafic et le tourisme via les données de format binaire du groupe d'experts du protocole de transport, génération 1 (TPEG1)

Partie 9: Événement trafic compact (TPEG1-TEC)

It focuses on the following requirements:

- ensuring travel safety for the driver;
- enabling the calculation of alternative routes;
- avoiding delays (e.g. traffic jams);
- warning the driver of obstructions on route;
- informing the driver of infrastructural problems (e.g. closed petrol stations, non-functioning emergency phones).

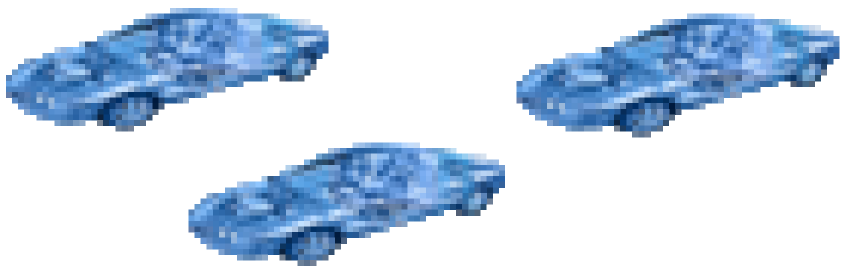

Data processing in the Automotive case

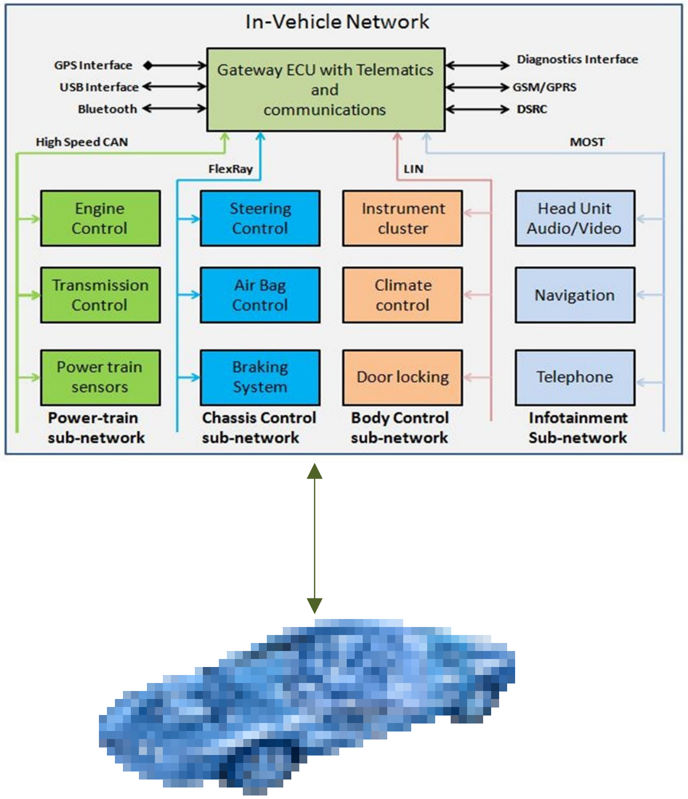

Real-time (μSec/mSec Level)

: Hard Real-time(< 10 ms)
: Semi Real-time(< 20 ms)
: Soft Real-time(< 100 ms)
: Non Real-time

On-Line (Sec Level)

Off-Line (hourly/weekly/ monthly Level)

Data processing in the Automotive case: possible examples

Scenario	Description	Output
Fault Diagnosis	Correlation between multiple sensor data	Vehicle Action (reporting failure)
Geo-fencing	Detection of enter/exit of the area	Vehicle Action (enter/exit warning)
Simple Event Detection	Event detection by sensor data of single vehicle	Event
Anti-theft Security	Security while parking	Vehicle Action (sending e-mail)

Scenario	Description	Output
Event Detection with multiple vehicles	Correlation between the same sensor data of multiple vehicles	Event
Complex Event Detection with multiple vehicles	Correlation between different sensor data of multiple vehicles	Event
Geo-fencing	Check number of vehicles in a specified area	Event

The Data Deluge effect upon IT infrastructure and Data Center architecture

- The Data Deluge indicate that enterprises will be dependent on 10 times the IT capacity of today, yet they will not have 10 times the budget to deliver.

- Current legacy datacenter architecture, in most cases, wasn't designed to meet the demands of a data-heavy era.

- Moreover, the average age of a data center, according to a recent IDC study, is nine years (Gartner suggests any data center more than seven years old is obsolete).

- Finally, power consumption and cooling demands, along with physical space costs represent first-order constraints in rethinking datacenter architecture.

Therefore, what can enterprises do to leverage their existing data center investments yet position themselves to be able to process an exploding volume of data?

Factors driving Data Centre design

Data Centre design is usually driven by 7 major factors.

	Details
Geographical Location	**Environmental risks, natural risks, human risks**
Facility Infrastructure	**Building type, security, cabling, fire protection, water detection, floor, control systems, facility management, etc.**
Electrical Power	**Power supply, UPS and Power Generators, indoor lights, etc.**
Cooling	**Topology of cooling systems, Temperature and humidity sensors, climate control, Hotspot Prevention measures, etc.**
Network	**Internal and external connectivity**
Services	**Services provided by the facility operator (security, building monitoring, building maintenance, etc.)**
Availability	**Ability to support continuous business operation**

The Availability factor in a Data Center design

Availability is implemented as one the of the following options:

Dual room

- In this case there is only one site, and the Computer Room is split into 2 rooms.
- In case of local outage (e.g. Fire, flooding, local explosion), there is a very high risk that both rooms and their equipments are impacted by the outage, triggering a break in availability of the IT services

Dual building

- In this case both rooms are separated by some distance (10th of meters or about 100 m)
- These rooms are in the same building , and the Data Center resilient infrastructures (Energy feeding and distribution, cooling, etc.) are built to support independently both rooms.
- In case of local outage, depending on how big is this outage, there is still a high risk of unavailability of the IT services

Dual sites

- In this case both rooms are in a different building in a different site, located at a distance of 10 to 20 km
- The data center facilities are independent and connected though dark fibers to ensure that both sites can operate in parallel. Clusters and storage are dispatched at each site, with replication of data to ensure minimal data loss.

Current data center topologies: the two-site topology (1/2)

- A two-site topology can operate in either active/passive or active/active mode.
- In an active/passive configuration, production workload is placed in the primary (active) site A and non-production workload, such as application development or testing, is placed in the secondary (passive) site B. Unidirectional Synchronous Disk Replication (SDR) is used to replicate the production data from the primary to the secondary site. The capacity of the secondary site is sized to support the production workload in case of primary site failure. Recovery at the secondary site involves restarting the mission critical business system and reconnecting end users.

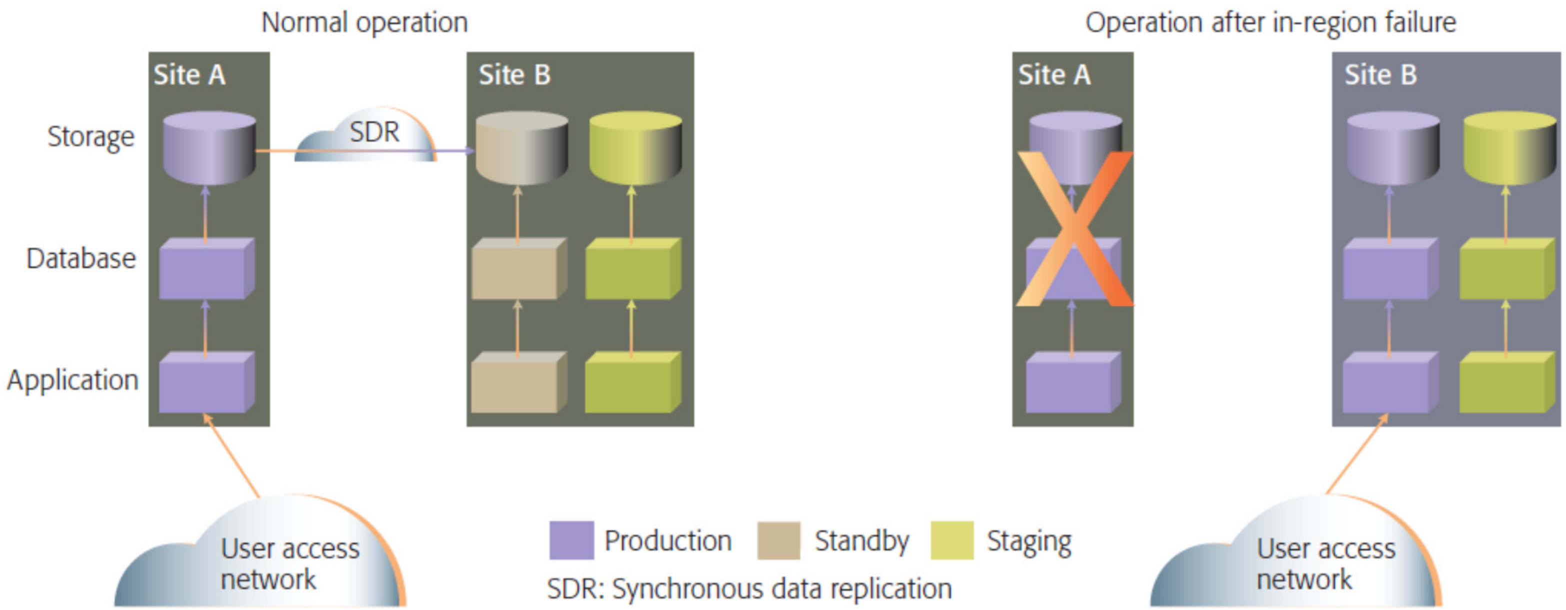

Current data center topologies: the two-site topology (2/2)

- In a two-site active/active configuration, load balancing in the application tier allows to direct transactions to either primary site A or secondary site B.
- Clustering in the database tier is accomplished with a shared database and a split mirror disk subsystem. Under normal operation, the primary-site database manager updates the database. Then, unidirectional synchronous disk replication is used to transmit a copy of the production data electronically from the primary to secondary site. Following a significant primary-site failure event, recovery at the secondary site involves restarting the database using the secondary-site database manager and reconnecting end users.

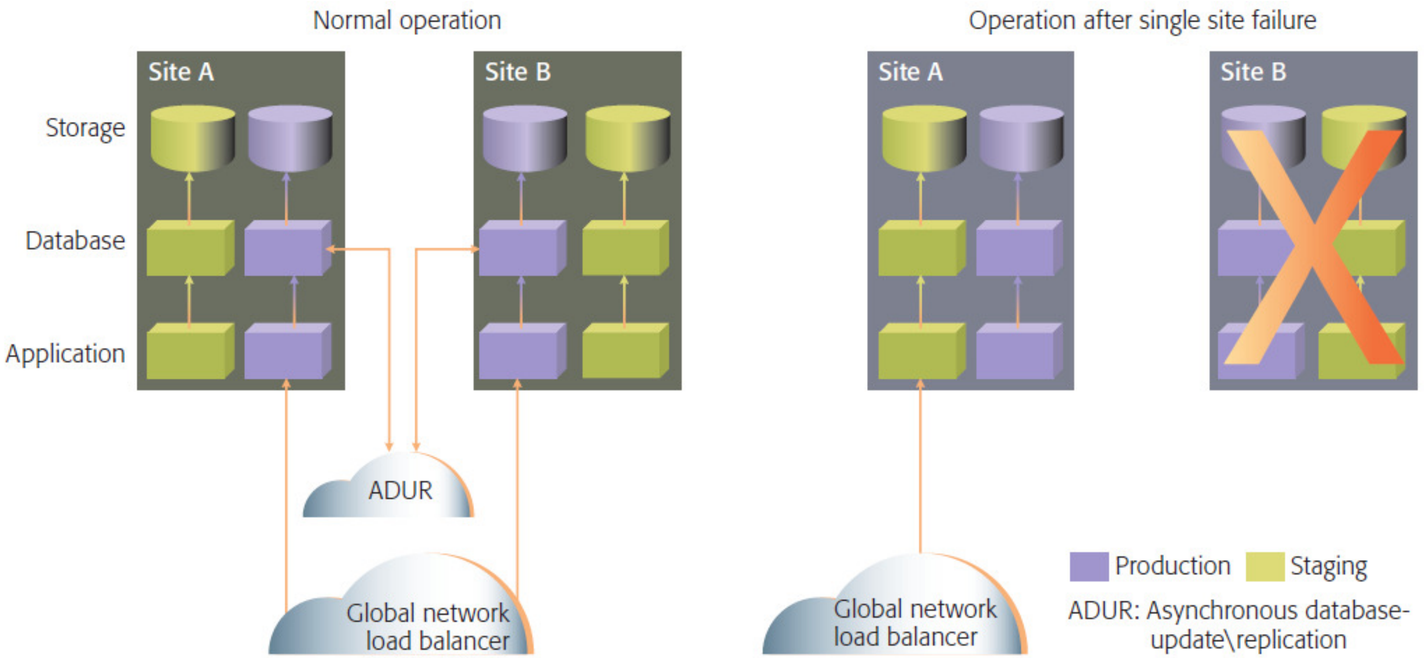

Reference architecture of a dual-site active-active configuration

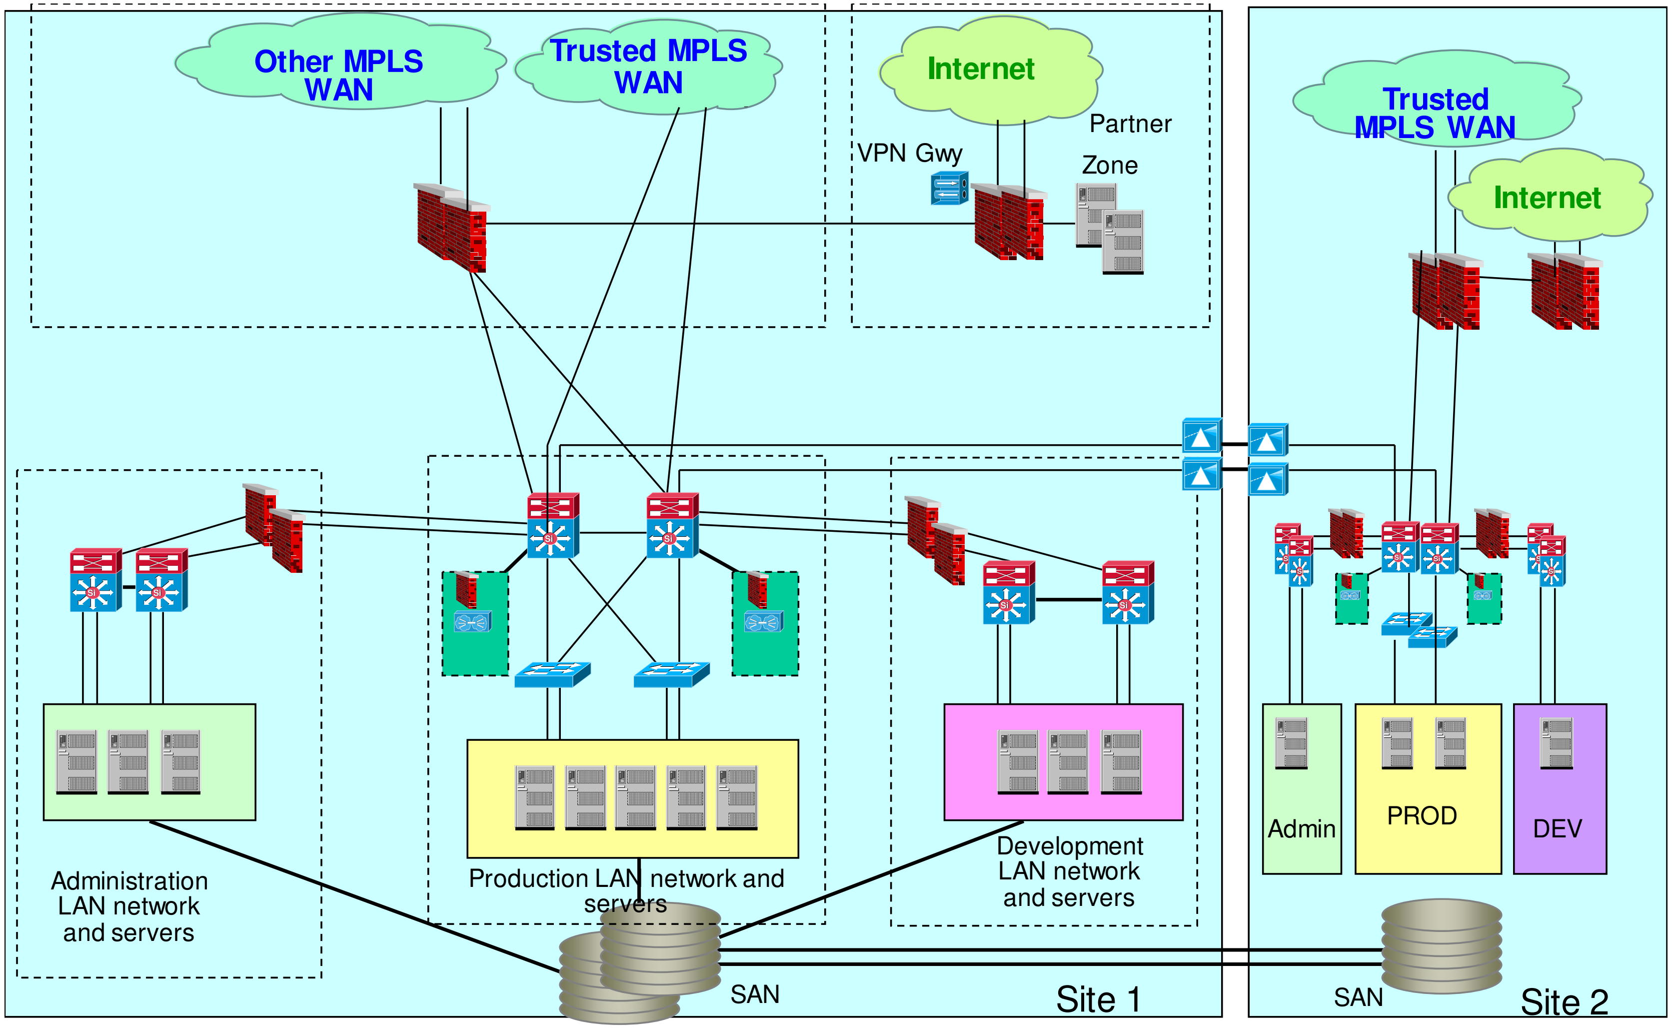

Current data center topologies: the three sites topology

- The three-site topology adds a third, out-of-region site to the two-site topology. It may be operated in either the active/active/passive or the all-active configuration.
- In the active/active/passive configuration, the two in-region sites A and B operate as a two-site active/active (shared) configuration. Under normal operation, the out-of-region site C acts as the passive site. Non-production workload is supported by the passive site. Database updates are sent asynchronously from production site B to the passive site C.
- The startup at site C involves restarting the mission critical business system and reconnecting end users.
- The three-site all-active configuration adds a third active site to the two-site active/active, shared-nothing configuration.

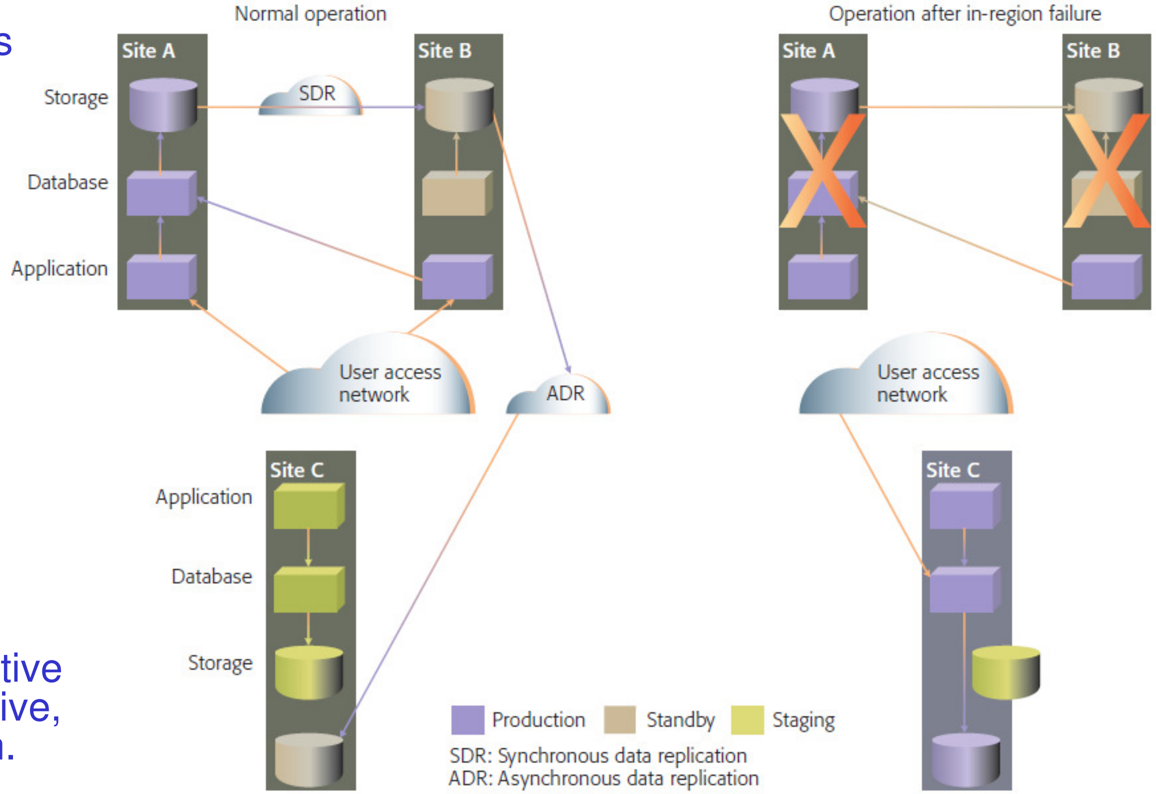

Comparison of data center topologies

Topology	RPO for each failure scenario			RTO for each failure scenario		
	Single-site failure	Regional disaster	Cascading failure	Single-site failure	Regional disaster	Cascading failure
Two-site active/standby	Minimal data loss	No protection	No protection	Recovery in hours	No recovery	No recovery
Two-site active/active	Minimal data loss	No protection	No protection	Recovery in minutes	No recovery	No recovery
Two-site shared-nothing	Minutes of data loss	Minutes of data loss	No protection	Recovery in minutes	Recovery in minutes	No recovery
Three-site active/active/ standby	Minutes of data loss	Minutes of data loss	No protection	Recovery in minutes	Recovery in hours	No recovery
Three-site all-active	Minutes of data loss	Minutes of data loss	No protection	Recovery in minutes	Recovery in minutes	No recovery
Four-site active/ active pairs	Minimal data loss	Minutes of data loss	Minimal data loss	Recovery in minutes	Recovery in hours	Recovery in minutes

Parameters in choosing a data center topology

FINANCIAL

- Cost effectiveness - hardware/software
- Cost effectiveness – site
- Cost effectiveness – personnel
- Cost effectiveness – network
- Cost effectiveness – transition
- Cost effectiveness – procurement
- Investment required for new facilities

ORGANIZATIONAL/CULTURAL

- Reuse of existing staff
- Political/geographical influence
- Cultural difference
- Primary Data Center close to business
- Transition impact on business operations

TECHNICAL

- Net latency
- Performances
- DR efficiency
- Operational efficiency
- Security
- Data centre load balancing
- IT management tools harmonization
- IT processes standardization and harmonization
- IT infrastructure sharing
- Technology standardization
- Exploit new technologies

Data Center resilience levels

	TIER I	TIER II	TIER III	TIER IV
Number of delivery paths	Only 1	Only 1	1 active 1 passive	2 active
Redundant components	N	N+1	N+1	2 (N+1) or S+S
Support space to raised floor ratio	20%	30%	80-90%	100%
Initial watts/ft²	20-30	40-50	40-60	50-80
Ultimate watts/ft²	20-30	40-50	100-150	150+
Raised floor height	12"	18"	30-36"	30-36"
Floor loading pounds/ft²	85	100	150	150+
Utility voltage	208, 480	208, 480	12-15kV	12-15kV
Months to implement	3	3 to 6	15 to 20	15 to 20
Year first deployed	1965	1970	1985	1995
Construction $/ft² raised floor*	$450	$600	$900	$1,100+
Annual IT downtime due to site	28.8 hrs	22.0 hrs	1.6 hrs	0.4 hrs
Site availability	99.671%	99.749%	99.982%	99.995%

*Excludes land and abnormal civil costs. Assumes minimum of 15,000 ft² of raised floor, architecturally plain one story building fitted out for the initial capacity, but with the backbone designed to reach the ultimate capacity with the installation of additional components. Make adjustments for NYC, Chicago, and other high cost areas.

© 2001 Uptime Institute

Resilience Level	Description
1	Basic data centre with no power protection and minimal resilience/redundancy of facility infrastructure.
2	Has UPS, but no generators
2+	Level 2, with the ability to perform concurrent maintenance on all mechanical and electrical systems without impacting operations
3	Has UPS and generators. Requires periodic shutdown for maintenance
3+	Level 3, with concurrent maintenance without site shutdown
3++	Level 3++ with automatic switch to alternate power source if UPS fails
4	High levels of resilience & redundancy, with two active UPS-protected power paths

IBM

The Compound Data Center paradigm

- Just as a molecular structure (compound) is comprised of several elements each of which can function independently but when bound together have new capabilities…, we can also imagine
- …a new data center paradigm with separate data centers, each one optimized for specific purposes, all together combined to optimize data processing and to provide the insights for decision making.

- We call such a **distributed and layered data center** paradigm "**Compound** (*) **Data Center**".

- Each element (Edge, Cloud, Enterprise) in the Compound Data Center has a specific role but when orchestrated together provide additional value.

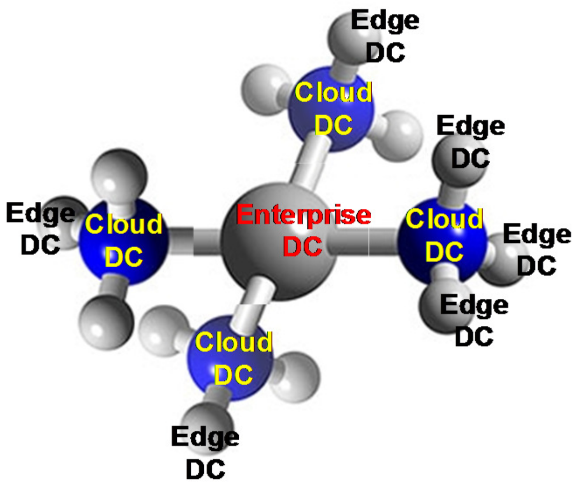

(*) "Any substance consisting of two or more different types of atoms (chemical elements) in a fixed proportion of its atoms can be termed a chemical compound (or just compound)" (source: https://en.wikipedia.org/wiki/Chemical_compound)

Infrastructure of Tomorrow: the general vision

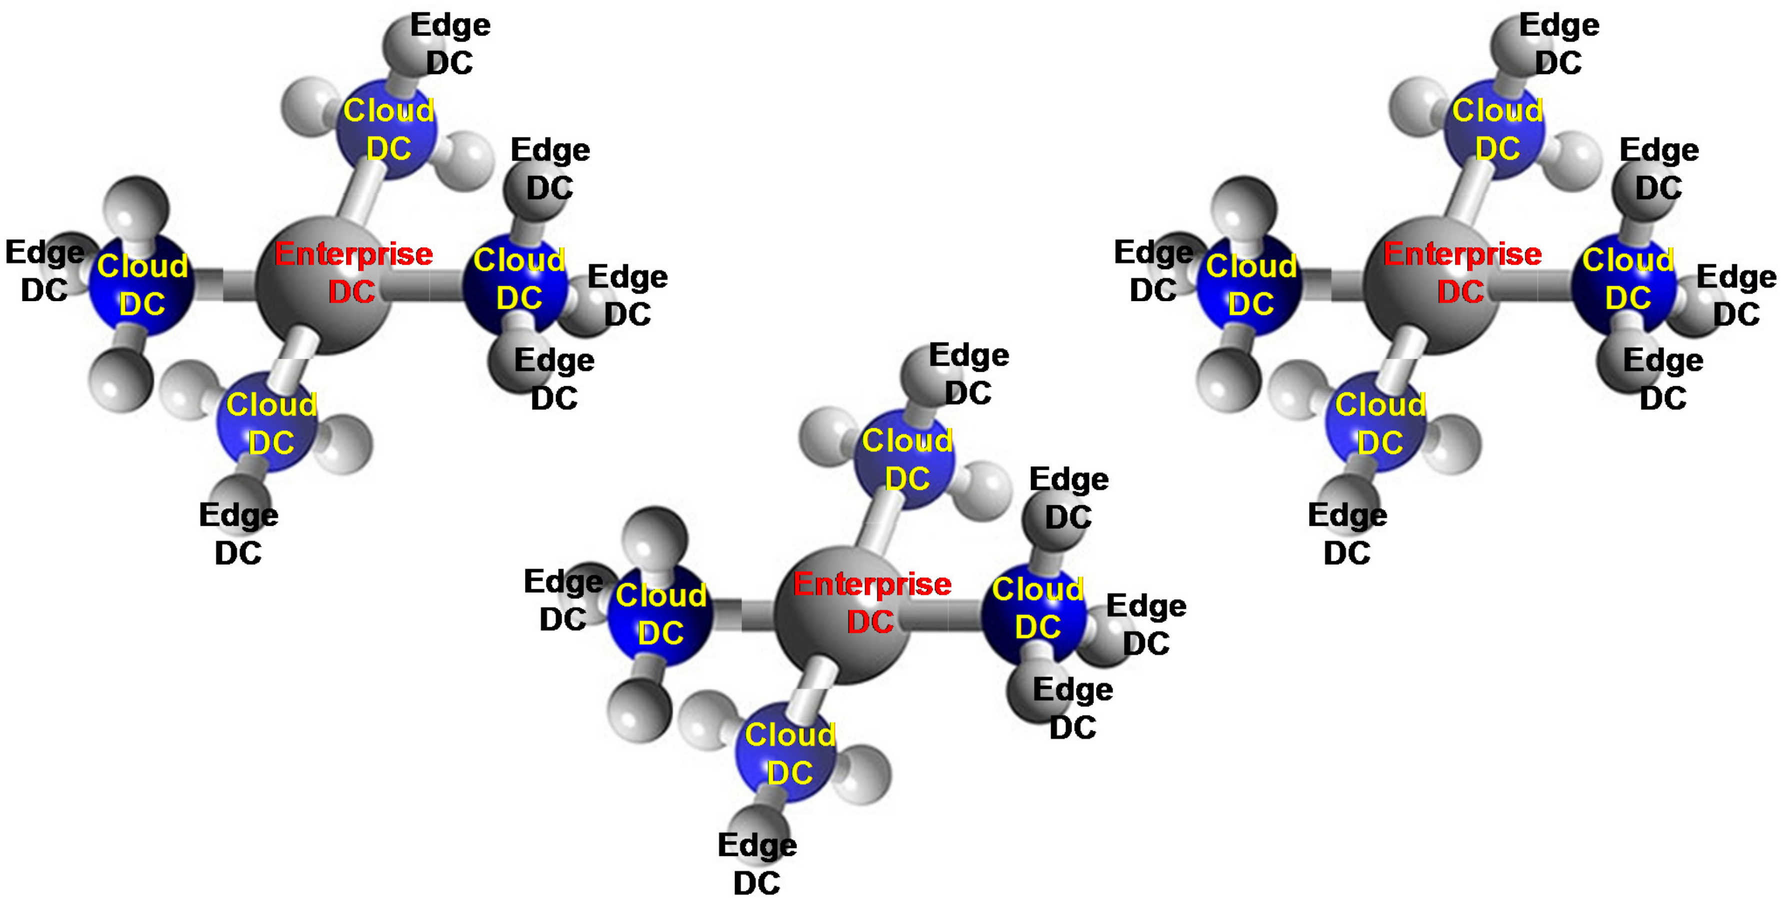

Outer layer: the Edge Data Center

- At the outermost layer, there will be **Edge Data Centers** (EDC**)** which will provide services for data collection (including real-time data streaming), noise filtering, formatting, distillation/reduction, aggregation, expansion, security, etc. for IoT devices in data center proximity.

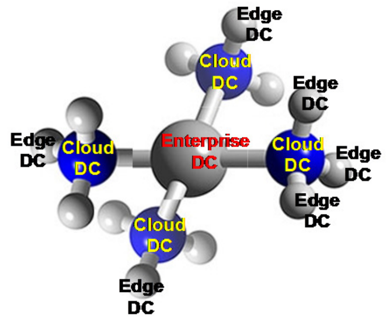

- Thus, the role of the Edge Data Center would be to process local data by harvesting and integrating the received data and discrete events for more comprehensive analysis.

- There are a number of reasons dictating the need for an Edge Data Center. For instance:
  - data affinity to the source of origin (e.g., sensor data) can help reduce latency.
  - data communications may either be impractical (e.g., large bandwidth required for massive amounts of data) or the transfer of data may violate local regulations or have privacy implications.
- Depending on specific data application and usage there could be a need to be one or more Edge Data Centers in a Compound Data Center.

Middle layer: the Cloud Data Center

- The middle layer, **Cloud Data Center**, will provide services to all edge data centers in its "service" range.
- The Cloud Data Center main mission is to derive "knowledge" from events coming from the Edge Data Centers.

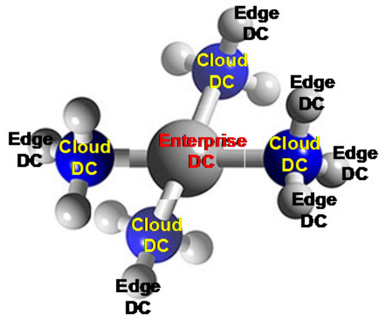

- In fact, event information will be augmented, integrated and correlated with data coming from various sources. Next, a semantic analysis, exploiting common vocabularies and topical ontologies, could create a common understanding of the data. Finally, analytics algorithms will extract actionable information (knowledge) to be used by higher-level applications and services.
- In the end, the Cloud Data Center will mainly provide large-scale analytics capabilities with associated scalable and dynamic computing power and storage. Moreover, the Cloud Data Center will make available deep learning services for data representation and abstraction.
- As at the outer Edge Data Centers, this layer could be comprised of multiple Cloud Data Centers depending on application and usage.

Inner layer: the Enterprise Data Center

- In the inner layer, we have the **Enterprise Data Center** running enterprise business applications which will use the knowledge from middle-layer Cloud Data Centers to execute the most appropriate actions.

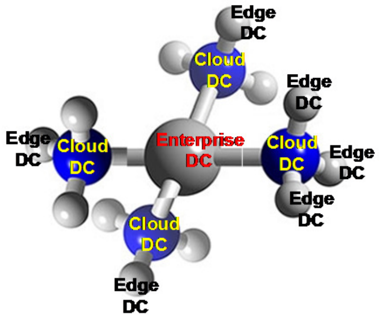

- These actions can contribute knowledge back to the Edge Data Center to create (or build on) a knowledge base and reference source for future situations.
- The Enterprise Data Center also directs the control and orchestration of all resources within the Compound Data Center as well as the management of the provided services.

The bound effect of the Compound Data Center paradigm

- Data is treated and transformed as it moves through the data processing chain.

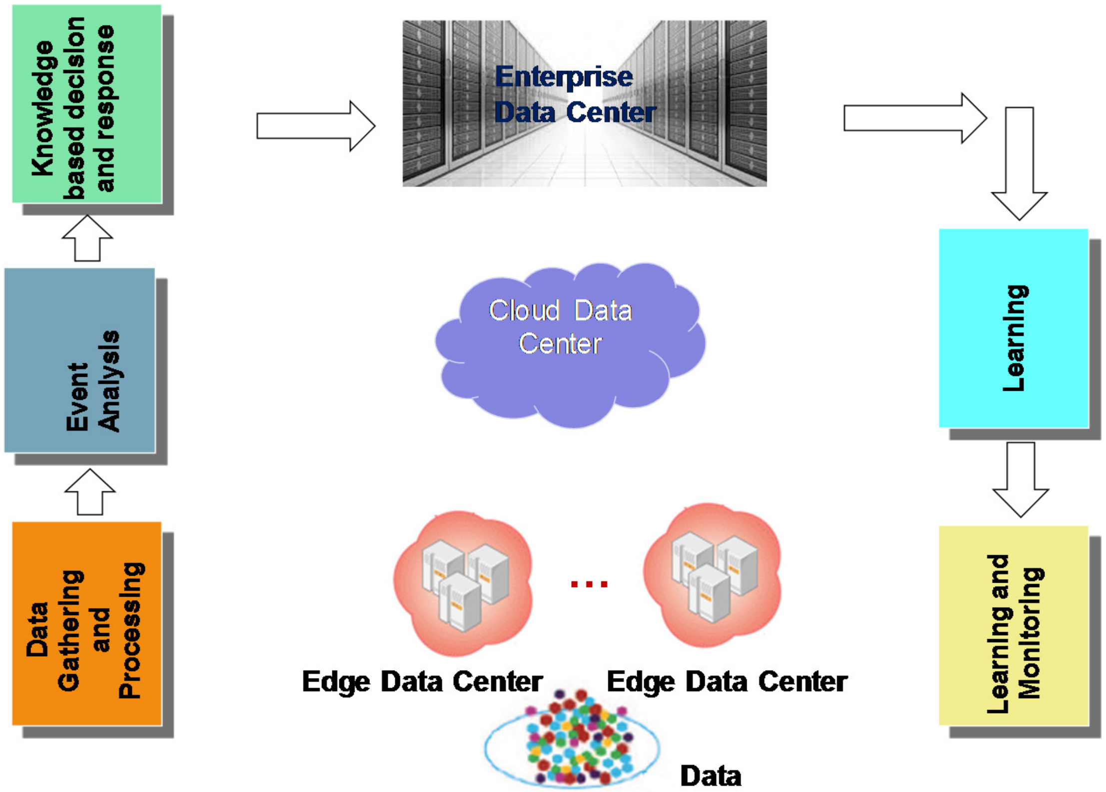

The final effect: from Data to Event, to Knowledge, to Action/Response.

The Automotive case: the predominant IT operating model

1.0 Develop Vision and Strategy	
2.0 New Vehicle Development	
3.0 Market and Sell Products and Services	
4.0 Build Vehicle and logistics	
5.0 Manage Customer Service	
6.0 Develop and Manage Human Capital	
7.0 Manage Information Technology	
8.0 Manage Financial Resources	
9.0 Acquire, Construct, and Manage Assets	
10.0 Manage Enterprise Risk, Compliance, Remediation and Resiliency	
11.0 Manage External Relationships	
12.0 Develop and Manage Business Capabilities	

(*)

automated through

Applications

operating upon

Centralized IT Infrastructure

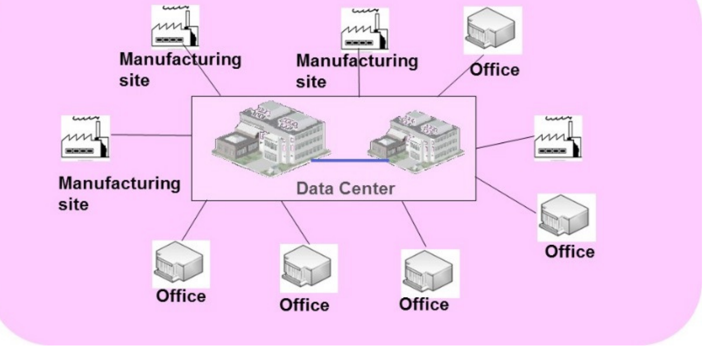

(*) source APQC, Automotive Process Classification Framework (PCF), v. 6.1.0

The Compound Data Center paradigm applied to the Automotive case: the idea

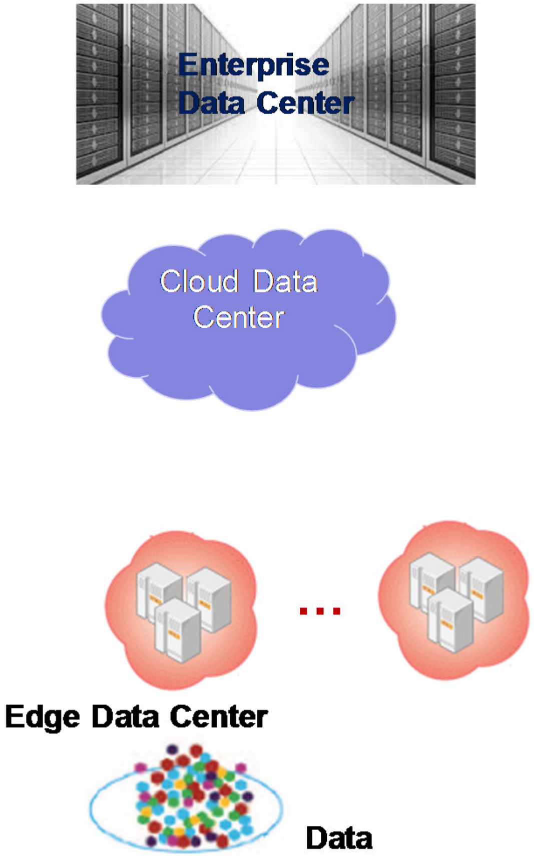

- **Systems of Record** (**SoR**) (enterprise business applications).
- **Systems of Insight** (**SoI**) (driver insights, vehicle insights, road infrastructure insights, etc.) mainly providing large-scale analytics capabilities.
- **Systems of Engagement** (**SoE**) offering “situation/ambient aware” services (e.g. location based services, context-related services, Intelligent Transport Systems ETSI TS 102 637-2/637-3 services, other highly personalized services, etc.) near real-time as well as information management services.
- **Data** (e.g. ISO ISO 22837:2009, vehicle probe data for wide area communications, etc.)

The Compound Data Center paradigm applied to the Automotive case: the consequence (an example)

1.0 Develop Vision and Strategy
2.0 New Vehicle Development
3.0 Market and Sell Products and Services
4.0 Build Vehicle and logistics
5.0 Manage Customer Service
6.0 Develop and Manage Human Capital
7.0 Manage Information Technology
8.0 Manage Financial Resources
9.0 Acquire, Construct, and Manage Assets
10.0 Manage Enterprise Risk, Compliance, Remediation and Resiliency
11.0 Manage External Relationships
12.0 Develop and Manage Business Capabilities

5.1 Develop post vehicle sale, customer care/customer service strategy (12635)

5.2 Plan and manage customer service operations (10379)

5.3 Perform after-sales installations and repairs i.e. vehicle service (12643)

5.4 Measure and evaluate customer service operations (10380)

5.5 Train and manage customer service work force (12651)

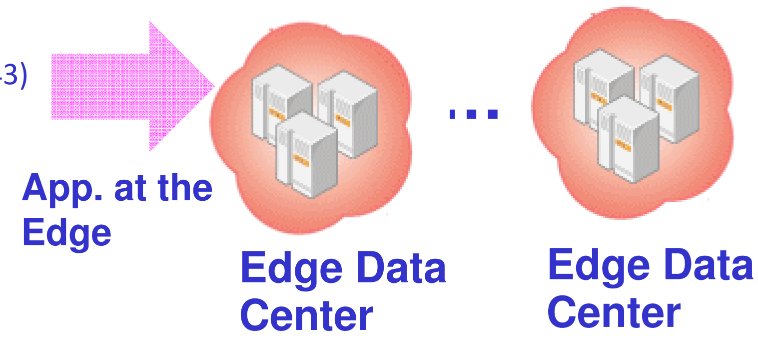

Cloud role in a Compound Data Center

- For mature organizations with a large investment in existing applications and data, it is unlikely that there will be a complete transition to cloud computing. This is especially true for critical Systems of Record that are at the hub of the Compound Data Center. Many core systems depend on existing applications and data. The cost and complexity of porting older applications often exceeds the benefit, leaving a nucleus of traditional data center technologies.

- Most likely, the next generation data center will see a mix of applications and data deployed in cloud and spread-out over the different layers of the Compound Data Center creating a **hybrid cloud** deployment model characterized by the adoption of technologies such as Application Programming Interfaces (APIs), Micro-Cloud, Containers, etc. with **Micro-Cloud** key to avoid massive data volume movements.

The Micro-Cloud technology

- In the Compound Data Center, a key design principle is to place analytics close to the data instead of moving data to analytics tasks.
- Thus, Micro-Cloud technology allows analytics tasks to exist in the form of micro-services (running in the Edge Data Centers) each executing separately and communicating via lightweight mechanisms (e.g., REST/HTTP).
- The Cloud Data Center acts as a central hub keeping the micro-services catalog for the Edge Data Centers.

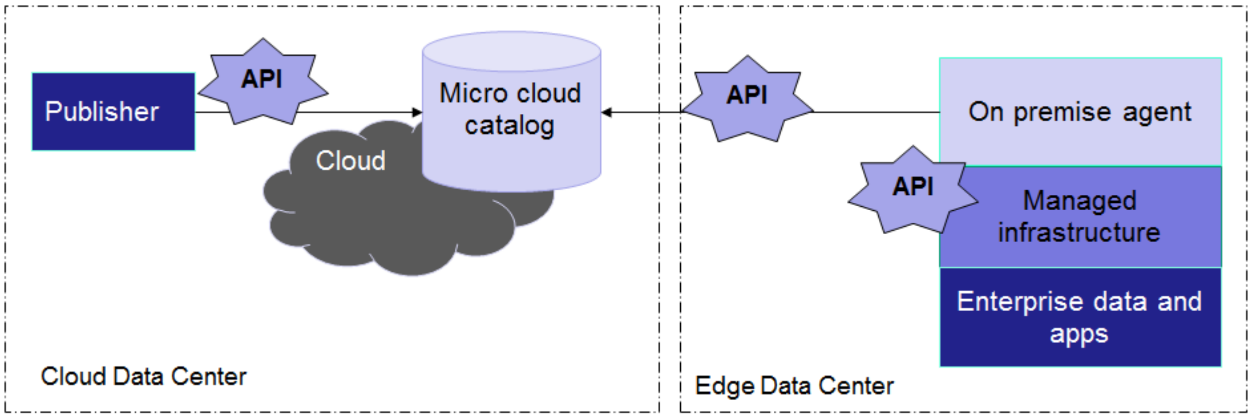

The Edge Micro-Cloud Appliance

- The Edge Data Center will host an appliance so to allow:
  - users to browse the public catalog and download required applications to the Edge.
  - application localization and deployment.
  - application execution in a safe and controlled environment.

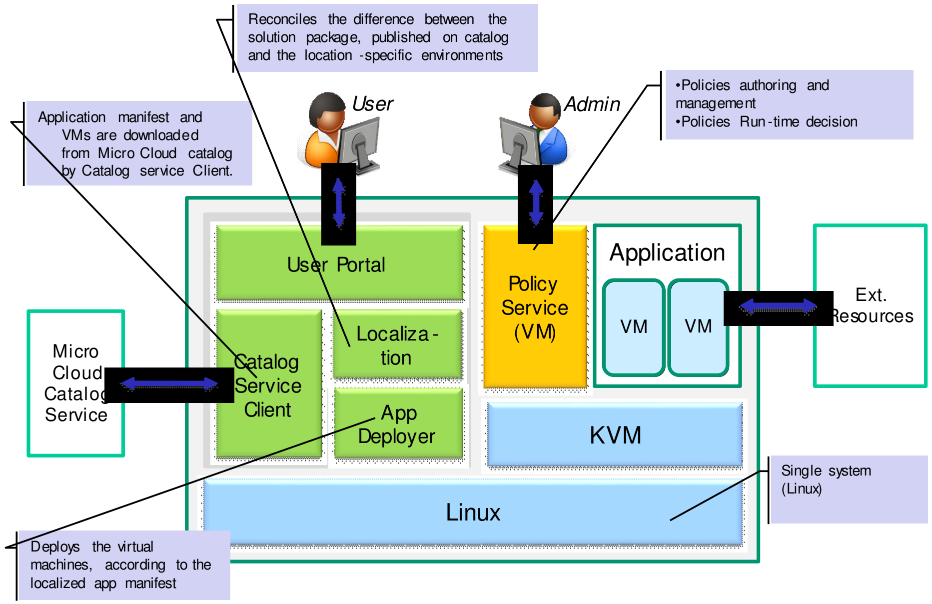

The Edge appliance can reduce management complexity, and allow on premise infrastructure at the Edge Data Centers to operate at costs and speeds comparable to those incurred in the cloud. The appliance can exploit solid-state technologies to optimize massive amounts of compute and networking power.

Compound Data Center physical overview

- From the physical standpoint, each data center in a compound should adopt a modular deployment method through the usage of pre-fabricated modules to realize scalability and a rapid delivery schedule.
- An Edge Data Center quite likely will be accommodated in an existing computer room (~20 m^2) or small data center (~150 m^2) where an IT infrastructure is present.
- The cloud data center is a typical large data center (~6000 m^2) usually established at a cloud provider location.
- The Enterprise data center (ranging from ~600 m^2 to ~6000 m^2) is the traditional legacy data center.

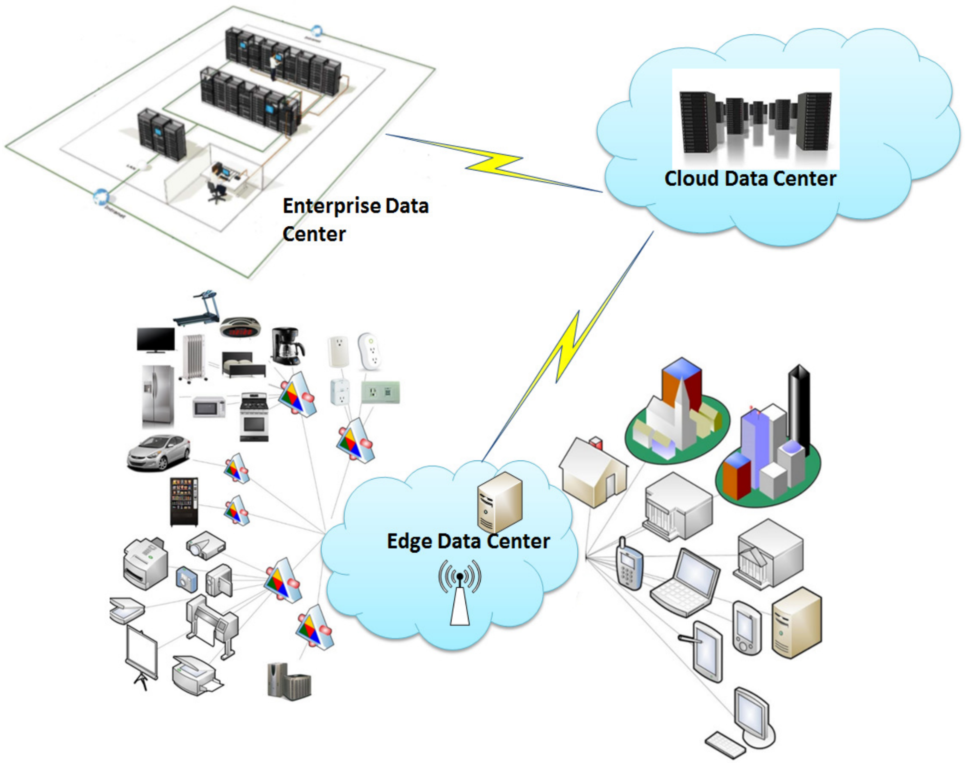

The interconnection in a Compound Data Center

- Data centers in a compound would be interconnected by means of an optical transport network (OTN).

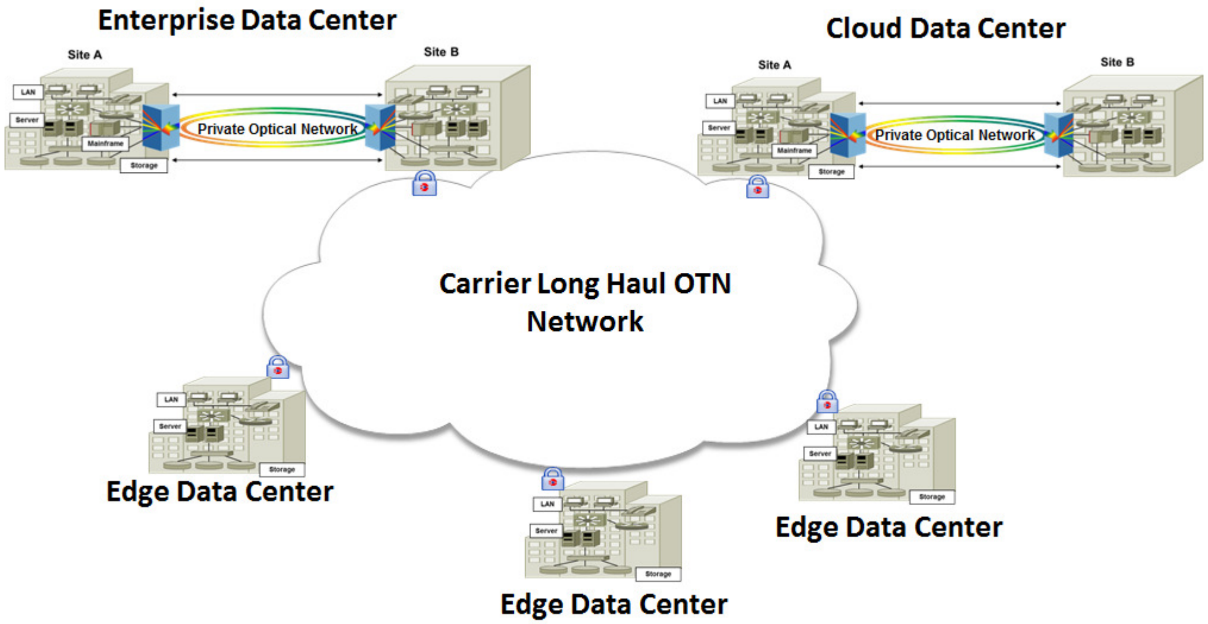

Requirements for a Compound Data Center

- There are many conflicting requirements when designing a distributed data center model. For example, data privacy might be of paramount importance but data protection (from hardware failures, site disaster and so on) could be just as important. And of course, one does not want to sacrifice performance with clients demanding almost immediate response time.

- Thus, the challenge in architecting a Compound Data Center is to achieve a balance between conflicting Functional and Non-Functional Requirements.

- We will focus next upon availability, security/protection and performance/scalability requirements for a Compound Data Center.

Availability in a Compound Data Center

- Although a higher availability tier may mean high availability levels, going for the highest level at each layer of the compound data center model is not necessarily the right decision.
- Factoring in the risk of downtime and potential business impact, having one or more Uptime Institute Tier II Edge Data Centers running in an active-active design could be a more cost effective alternative. In this way, a 99.5 (~2 days of unplanned downtime per year) availability percentage for the Edge infrastructure components (server, storage, network equipment) could be considered acceptable. A further improvement may come from the “software defined” management of the Compound Data Center to support load balancing of connections and applications across the existing Edge Data Centers as well as data replication to help avoid service disruptions.
- In the Cloud Data Center and the Enterprise Data Center elements of a Compound Data Center, infrastructure availability usually starts at a 99.8 (Uptime Institute Tier III, Tier IV) level (~18 hours of unplanned downtime per year).

Infrastructure protection in a Compound Data Center

- Although a distributed data center model looks to provide an efficient means of conducting extraordinarily large-scale data processing computations, numerous security issues must be taken into consideration once the model is implemented.
- In order to secure a compound data center, organizations can no longer depend on a traditional security approach that focuses on protection at the network level. Once the network is breached, hackers can easily access systems and data within the compromised network. In the compound, instead, there is the need to secure all infrastructure and data within its perimeter.
- Infrastructure protection technologies are typically segmented based on the component to protect. As also Gartner indicates "…*no single safeguard will protect against all possible attacks…*" and "…*Technology or services alone cannot provide effective infrastructure protection. Effective processes, as well as adequate deployment and operations staffing, are also required*". If these considerations apply to a traditional data center model, no doubt they are also valid for a compound data center.

IBM

Data protection in a Compound Data Center

- In a compound data center, data protection and security mechanisms should be in place when:
  - Data is at rest (in a device, in a storage system, in a server local disk).
  - Data is in transit along a network link.
  - Data is in use during a processing task.
  - Data is aggregated.
- In addition, there is the need to protect not only the data, but also the metadata.
- Over the last few years, several data protection technologies have been introduced. Among them encryption and tokenization have seen a strong adoption rate due also to the rapid diffusion of Bitcoin digital currency and the associated blockchain technology.
- We believe that in a compound data center, by means of the **blockchain** database and smart contracts among the parties, data processing tasks at each layer could securely communicate without relying on a central security system.

Performance and Scalability in a Compound Data Center

- Numerous design choices and solutions can be adopted so to have the Compound Data Center:
  - Executing workloads within a defined performance profile.
  - Handling increased processing volumes in the future if required.

- Some existing solutions, that represent a starting point to make decisions when dealing with performance and scalability for a compound data center, are:
  - Scale-out" and "scale-up" which are two traditional ways to provide more resources when necessary.
  - "Multi-processing/multi-threading" with applications running upon a server supporting parallel threads execution.
  - "Workload optimization" where a workload is matched to a specific configuration of server, middleware, storage and network components. This is the case of the Edge Micro-Cloud appliance.

Conclusions

- As Gartner's Analyst David J. Cappuccio expresses in the "The Future of the Data Center in the Cloud Era" June 2015 research note "…*Different application types will reside where the delivery model best supports client expectations, risk, compliance, service continuity and regulatory issues. This model moves IT away from the traditional, on-premises, full-control model toward a distributed computing model…*".
- In fact, the Compound Data Center enables small local data centers to serve Systems of Engagement workloads in an efficient way without excessive network latency penalties. Moreover, they avoid the need to send data across a larger network thus reducing the risk of compromising privacy.
- The Cloud Data Center is the place where massive computing operations can occur, and is an alternative to the large investments that companies have spent when working with legacy data centers.
- The Enterprise Data Center allows companies to keep in-house the critical Systems of Record where a cloud solution is not appropriate.
- The software-defined model applied to the Compound Data Center enriches the overall architecture. It improves automation and flexibility in order to better sustain business agility and dynamic growth.