Scalability: From Hundreds to Thousands Architecture + Process

Richard L Johnson
Principal Technical Architect
AT&T Global Technology Operations
Service Continuity and Disaster Recovery

Why Do We Need To Talk About Scalability and Architecture?

Because Zerto is no longer just a Disaster Recovery tool.

Because Zerto has matured into an Enterprise Resiliency Platform that carries with it some pretty profound impacts to your organization.

And because your Zerto environment is getting bigger every day – from tens to hundreds, or from hundreds to thousands of VMs, scaling your environment is critical.

Understanding those impacts means:

- placing them in an enterprise context
- understanding how Zerto relates to the overall Enterprise
- understand how Zerto needs to operate, both now and in the future

AT&T

AT&T

“More IT Projects fail from poor architecture than from poor planning”*

* Poor planning is a close second, however.
  But without an architecture, what is going to be planned?

Example - The Minimum Viable Product

High Level Project Plan:

1) Ideation
2) Lab Test
3) Proof of Concept
   - Use Cases
     - Functional requirements
     - Non-functional requirements
4) Limited Production Rollout
   - Business Case & Budget
   - MVP
5) Rollout
6) Iterate

The MIAs:

- *Where's the architecture?*
- *Where's the operational plan?*
- *Where's the training?*
- *Where's the communications plan?*

AT&T

What An Architecture Does:

Provides a requirements-based framework that supports

- Innovation,
- Flexibility
- Growth

Balances needs of different constituencies

- Business Client
- Developers
- Operations
- Security

Provides a contract between those communities for success

- Planning
- Building
- Running the application

"Why do we need architecture?

- *To clarify, define, and ensure meeting the client expectation*
- *To align the business function with the technology*
- *But mostly, to avoid obvious violations of common sense."*

CIO of Fortune 100 Company
(Quote Allowed Anonymously)

AT&T

Types of Scaling

Vertical – Add more RAM, Faster CPU, more disk

- "Armstrong Method" e.g. brute force
- Typically expensive and static

Horizontal – Add more instances

- Works for front end, not so much for DB: Clustering, Sharding
- The Myth of Auto-Scaling

Algorithmic – (Re)use more efficient programming elements

- Catalog & repository of reusable components & patterns
- Requires better trained & experienced developers

Procedural – Better organize process flows

- Understanding requirements in context: business knowledge
- Deconstruction – the Architects Secret Weapon
- Star vs Continual processes

Operational – Automation and Tenant Management

- To scale, ultimately IT must surrender control to the tenant
- And IT must provide systems that allow that to succeed.

When we talk about scalability there are at least 5 major forms of scalability, and all of them are valid within the own contexts.

Understanding – or even having – a system reference architecture enables understanding where, what, and how to scale.

AT&T

Understand Enterprise Architecture Patterns

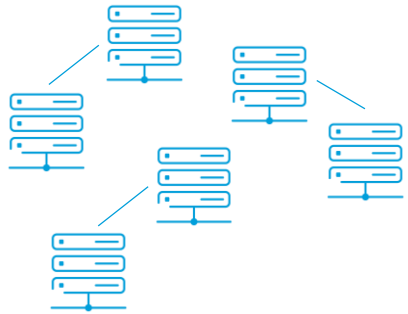

Geo Clusters – OK for DR, OK for DB, not so good for migration and growth.

- Management is by clustered pair.
- Capacity model p= (n+s)/2 Cost model is: c=2(n+s) + l
- Paired sites restricts flexibility.

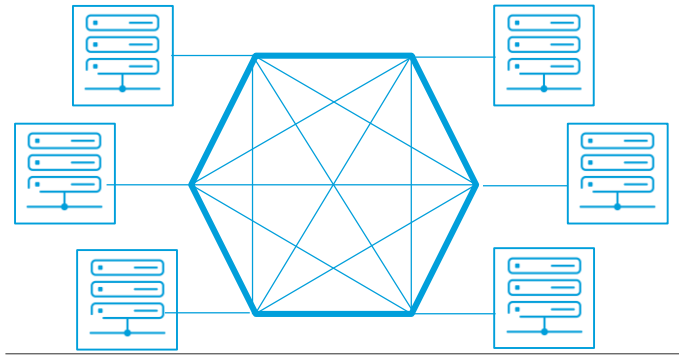

Full Mesh – Any to Any, with great flexibility comes with great complexity and cost.

- Compute resources and storage are in the edge nodes.
- Capacity is still p = (n+s)/2. Cost is c= 2(n+s) + l
- Networking connectivity and node management is in the mesh.

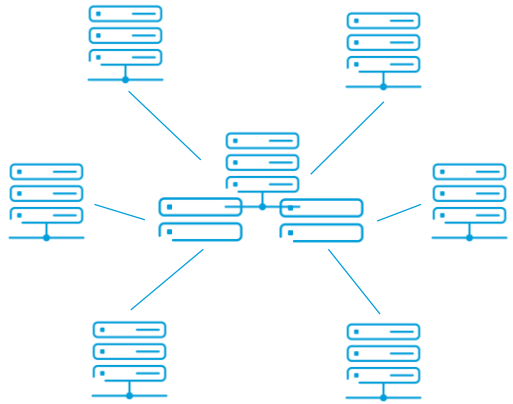

Hub and Spoke – Enables lower cost oversubscription protection model for compute. Still requires full storage

- Complex compute recovery model
- Capacity is p= (n-1) +(s/2) Cost Model is c=(n+1) + 2S + l
- This is the Public Cloud model.

Variables c = Cost p = Capacity n = Nodes s= Storage

AT&T

Planning For Big –

VRA Addressing Scheme /16 = 64K Hosts 32K allocated, 32K reserved

Data Center	Max Hosts	/24 needed	Subnet Start	Usable range for Zerto	Subnet Stop
	12,480	52	10.122.0.0	.10 - .250	10.122.51.255
	480	2	10.122.52.0	.10 - .250	10.122.53.255
	480	2	10.122.54.0	.10 - .250	10.122.55.255
	1,920	8	10.122.56.0	.10 - .250	10.122.63.255
	1,920	8	10.122.66.0	.10 - .250	10.122.73.255
	960	4	10.122.76.0	.10 - .250	10.122.79.255
	1,920	8	10.122.80.0	.10 - .250	10.122.87.255
	960	4	10.122.88.0	.10 - .250	10.122.91.255
	1,920	8	10.122.92.0	.10 - .250	10.122.99.255
	960	4	10.122.100.0	.10 - .250	10.122.103.255
	1,440	6	10.122.104.0	.10 - .250	10.122.109.255
	1,440	6	10.122.110.0	.10 - .250	10.122.115.255
	1,440	6	10.122.116.0	.10 - .250	10.122.121.255
	1,440	6	10.122.122.0	.10 - .250	10.122.127.255
	960	4	10.122.64.0	.10 - .250	10.122.65.255

AT&T

RASCI Chart All Your Processes

R = Responsible A = Accountable S = Supports C = Consults I = Informs

Business Process Zerto Replication Manager Implementation and Support	SCDR	VMW System Hosting	TCAM	System Deployment	Tools Team	ARC	Appl Support	EBS	PDBA - Ops	Appl DBA	Rationale
Production Environment - Virtual Replication Appliance (on each ESX host VM at each site)											
Install Zerto license and VRA at Site 1 on ESX hosts where replication is to be done - (ZVM installed at each site)	I	R/A									
Install Zerto VRA license for ESX hosts at Site 2	I	R/A									
Connect/Pair Site 1 & Site 2	I	R/A									
Tuning of VRAs (memory and cpus)	I	R/A									
Training of User Group (ARCs)											
Operational and Process Documentation	C	R/A									
Virtual Protection Group (VPG)											
Determine application eligibility for Zerto	C		R/A			C/I	C/I		C/I	C/I	
Determine affinity groupings (which applications VMs/database VMs to be protected)	C/I	S	R/A			C/I	C/I				Requires knowledge of the appl
Create and configure Virtual Projection Groups (VPG)	C/I	S	R/A			C/I	C/I				Requires knowledge of the appl
Pre-seeding (Instantiation) of initial Site 2 VMs	C/I	S	R/A			C/I	C/I				
Operational - Move/Commit and Rollback of VPGs											
Move VPG (group of VMs) from one site to another	C/I	S				R/A	R		C/I	C/I	ARCs currently do reciprocal failovers
Failover VPG one site to another	C/I	S				R/A	R		C/I	C/I	ARCs currently do reciprocal failovers
Commit or roll-back a move/fail-over	C/I	S				R/A	R	I	C/I	C/I	ARCs currently do reciprocal failovers
DR Exercise											
Pre-production fail-over testing	C/I	S				R/A	R		C/I	C/I	ARCs currently manage pre-prod exercices
Ongoing fail-over exercising of VPG one site to another	C/I	S				R/A	R		C/I	C/I	ARCs currently manage reciprocal exercices
Commit or roll-back from a fail-over test (if applicable)	C/I	S				R/A	R	I	C/I	C/I	ARCs currently manage reciprocal exercices
Validation Application Recovered	I	I				R/A	R		C/I	C/I	
Validation DB Recovered	I	I				A	C		C/I	R	

AT&T

Leverage APIs

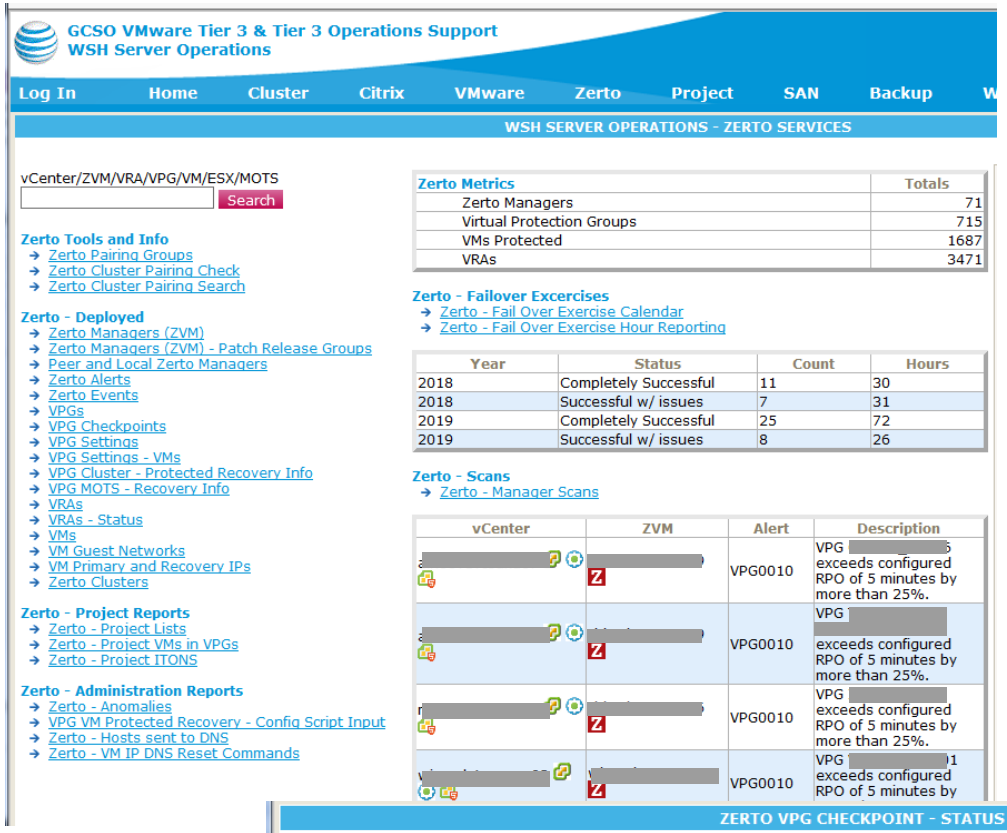

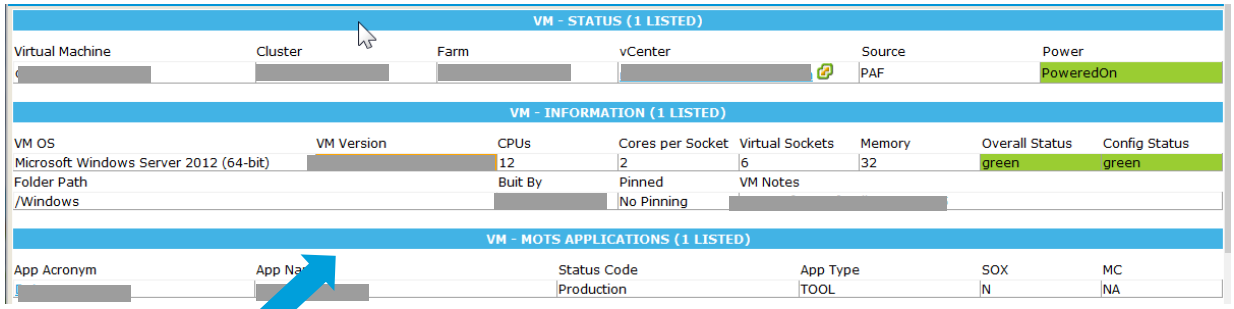

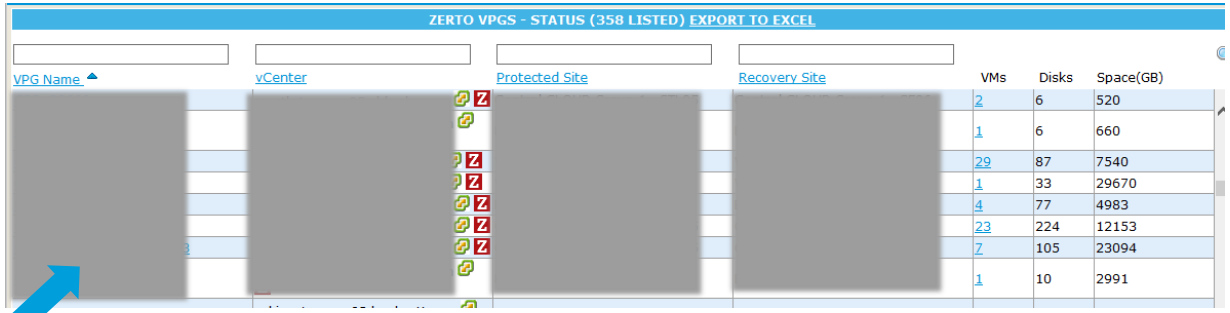

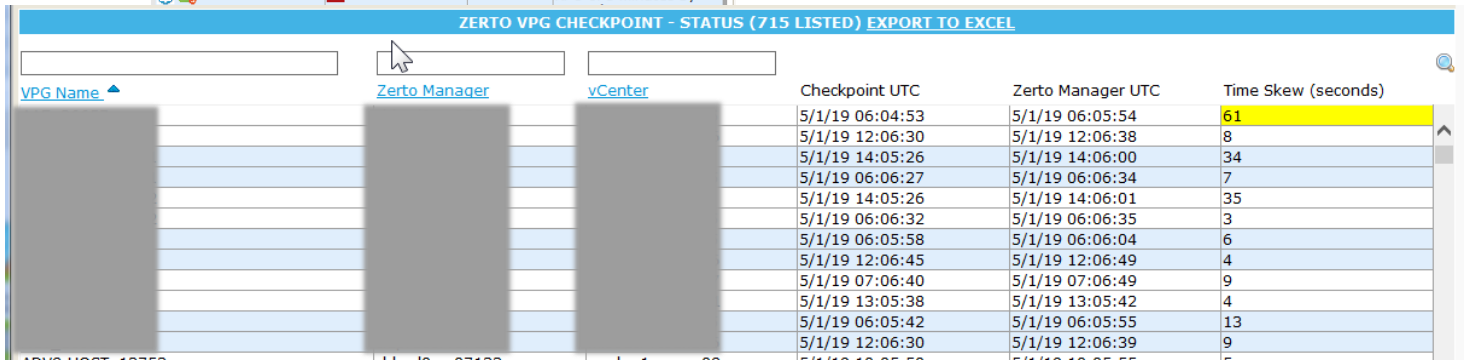

- We used both AT&T and Zerto APIs
- Drill-downfrom VPG to VM config detail
- All lists export to Excel

AT&T

Leadership Architecture Scorecard

Element	Completeness (0-5)	Maturity (0-5)	Communications (0-5)	Execution & Process (O-5)	Total
Vision					
Mission					
Strategy					
Goals					
Objectives					
Deliverables					
Resources					
Plan					
Governance					
Guiding Principles					
Totals:					

AT&T

And In Conclusion…

- ✓ Have an Architecture
- ✓ Communicate it, use it, refine it
  - Client
  - Developers
  - Operations
- ✓ Plan for Big
- ✓ Understand and Document
  - Processes
  - Roles
  - Responsibilities
- ✓ Leverage Available Tools & APIs
  - http://s3.amazonaws.com/zertodownload_docs/Latest/Zerto%20Scale%20and%20Benchmarking%20Guidelines.pdf
- ✓ Have a scorecard

Complexity is the enemy of scale.

Successful scaling requires simplification

AT&T

Richard Johnson, Principal Technical Architect for Service Continuity and Disaster Recovery at AT&T has been a DR practitioner since 2003. Prior to that he was an Enterprise Network Architect at AT&T.

He has also worked in enterprise network, architecture and management roles for GE Capital, Starbucks, Premera Blue Cross, Weyerhaeuser, and Pacific Telesis.

He has been working with Zerto since 2013.

AT&T