PURESTORAGE®

Keep Your Data Secure! Strategies for Business Resiliency

강신우

Senior Systems Engineer
PURE STORAGE KOREA

sean.kang@purestorage

Nov 2024

데이터가 기업의 가장 중요한 지원

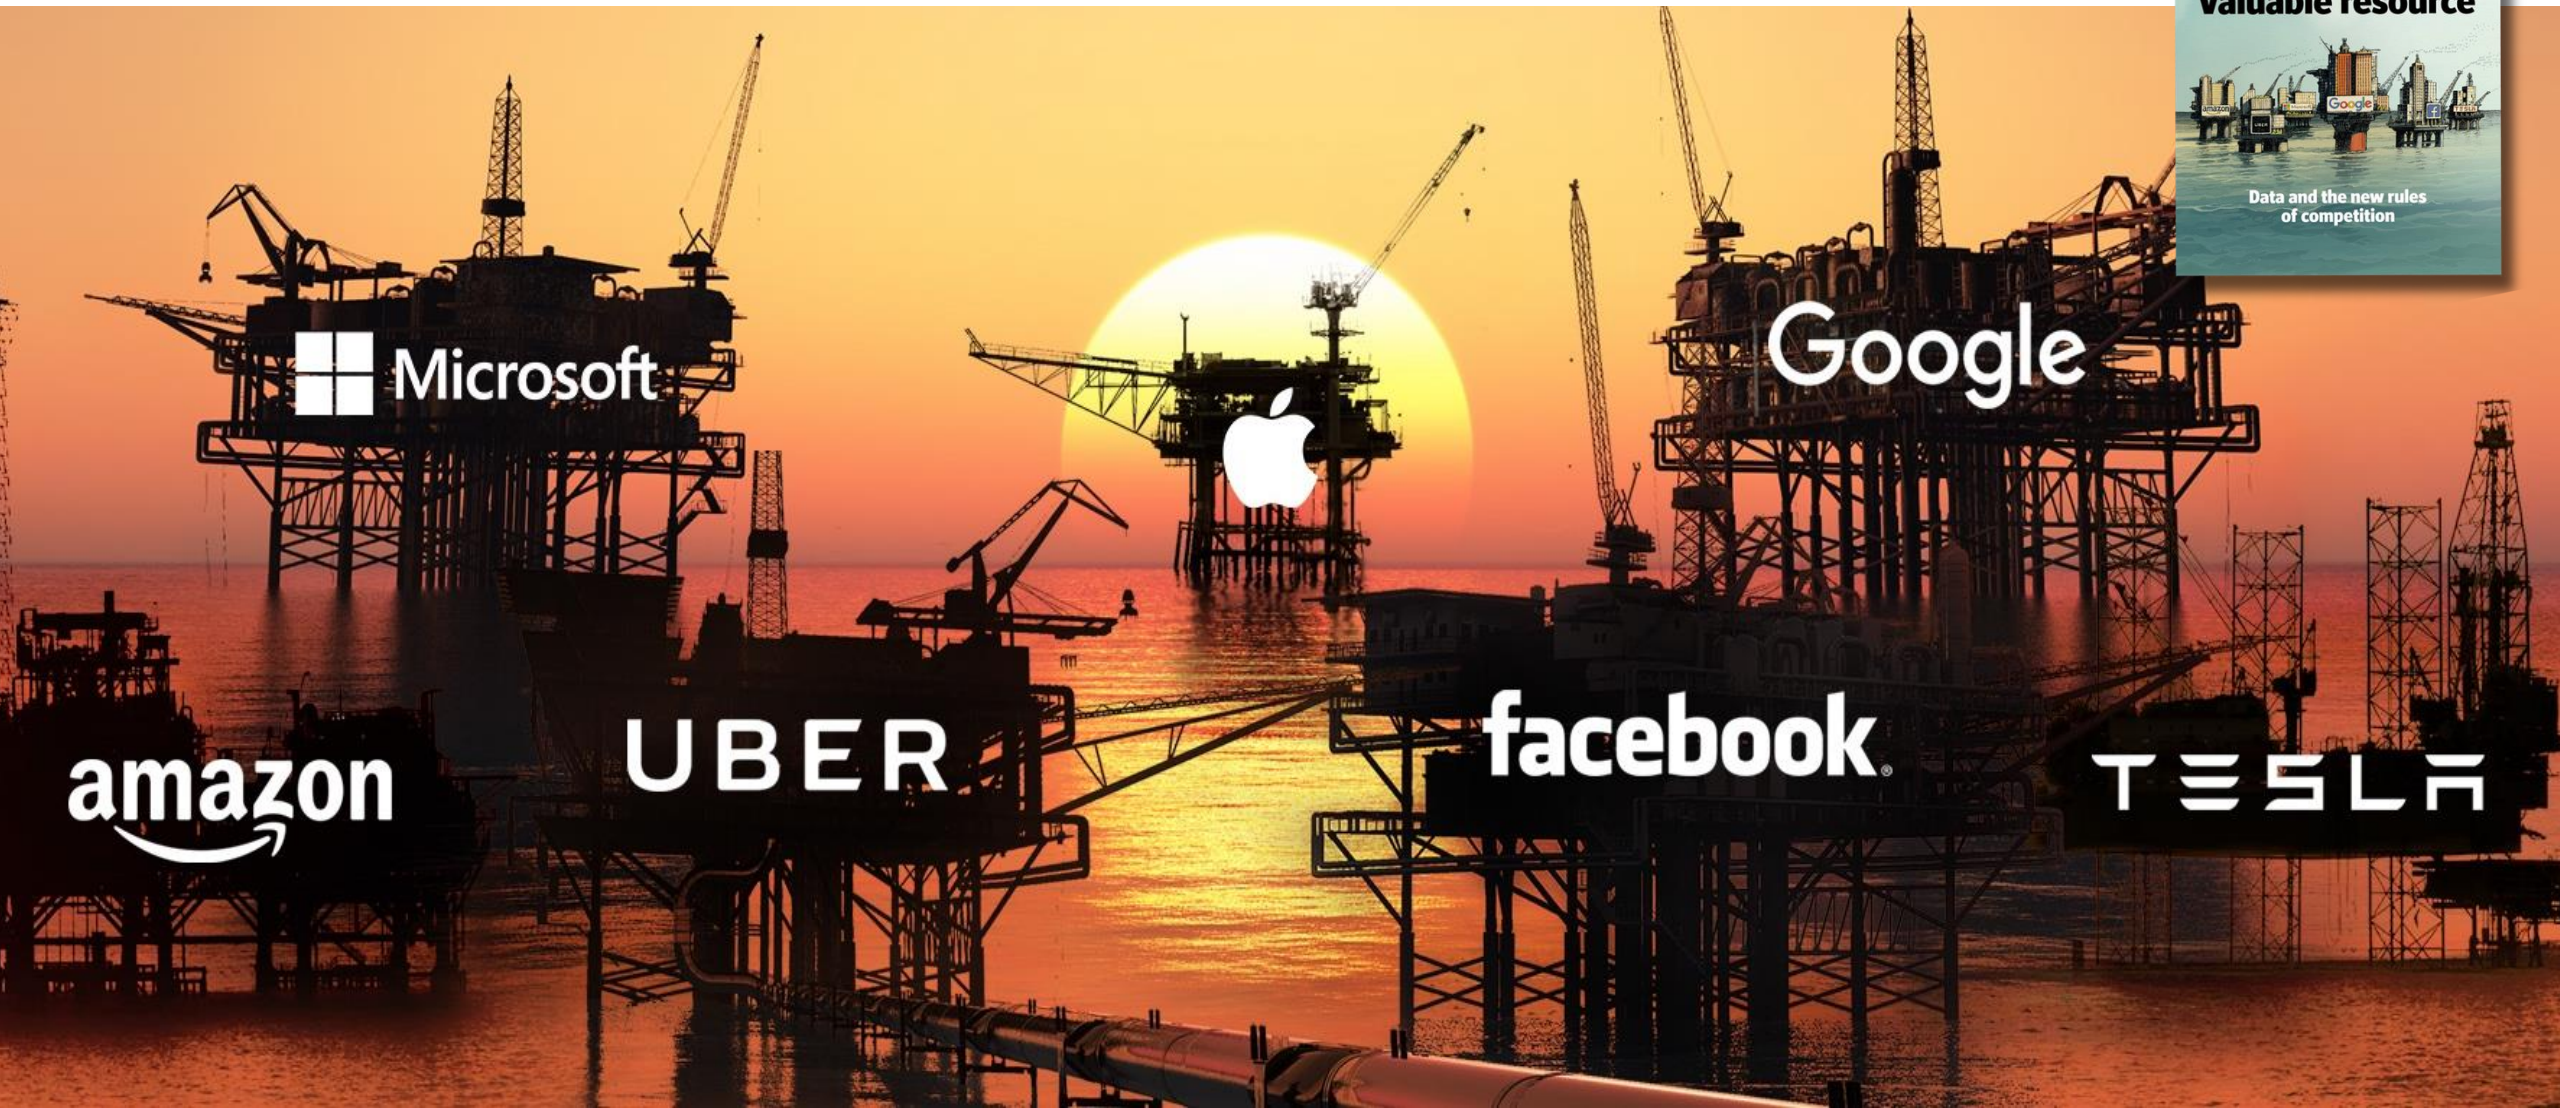

Re-Think Business Continuity

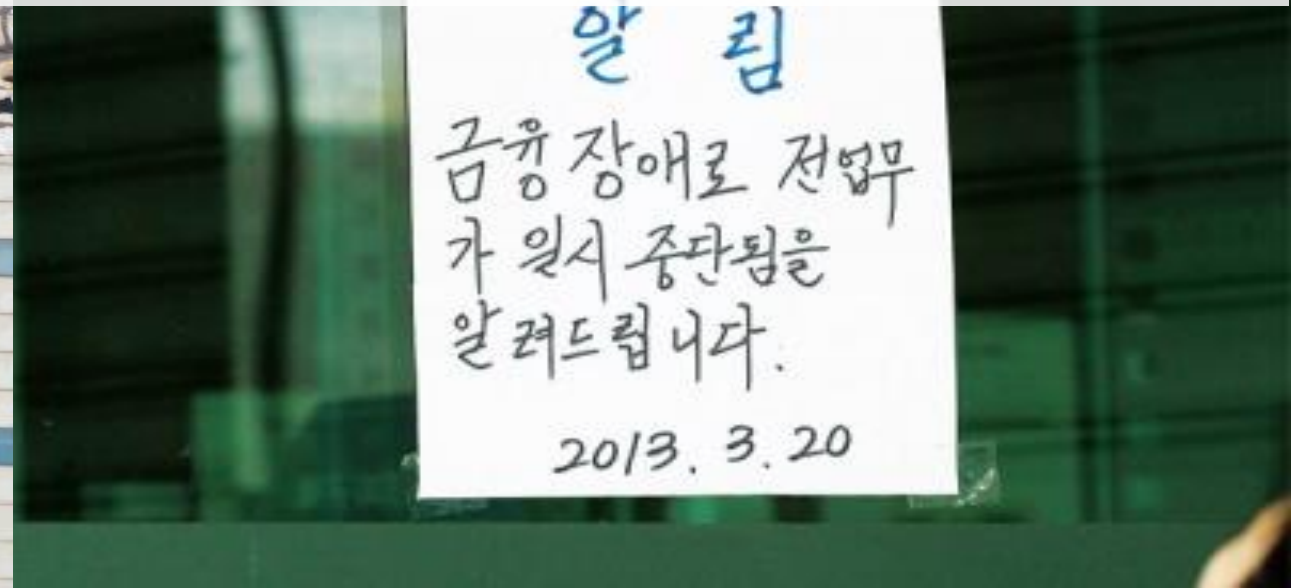

데이터 관점에서의 장애 유형

물리적 장애 [Physical Failure]
: 하드웨어 또는 전원, 데이터센터와 같은 물리적인 장비의 문제로 발생되는 장애

- 자연재해 - 지진, 폭우, 지진, 화재
- 하드웨어 노후화로 인한 실패
- 전기 / 네트워크 / 스토리지 장애

→ 하드웨어 이중화를 통한 무중단 운영

→ 백업을 이용한 데이터 손실 방지

재해 복구 RPO/RTO 수준 정의

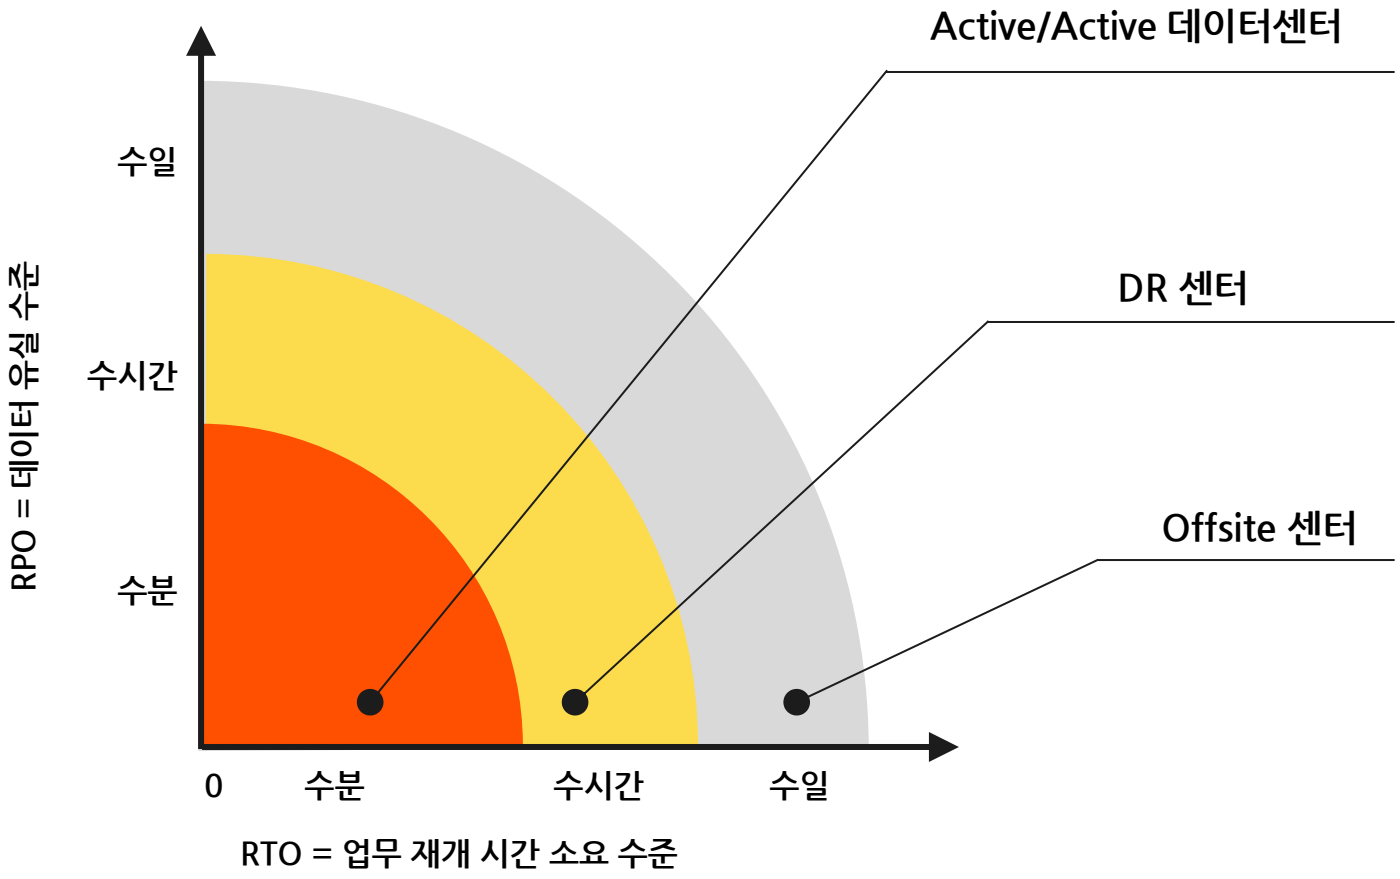

CONTINUOUS AVAILABILITY
- 서비스 무중단 (Zero Downtime)
- 최소중단(Near Zero Downtime)

DISASTER RECOVERY
- 중단 → 복구 → 서비스 재개
- Hour ~ Day

BACKUP & RECOVERY
- 중단 → 복구 → 서비스 재개
- Days

데이터 관점에서의 장애 유형

논리적 장애 [Logical Failure]
: 애플리케이션 또는 운영체제 상의
데이터 삭제 및 손상으로 발생되는 장애

- Human Fault
- 운영자 / DBA 에 의한 작업 실수
- 사이버 공격으로 인한 데이터 손실

→ 백업을 통한 데이터 복구

재래식 백업은 비즈니스를 보호할 수 있나요?

DCAC

If You Thought Database Restores Were Slow, Try Restoring From an ~~EMC Data Domain~~

By: Denny Cherry

Published On: July 20, 2015

Recently I did something which I haven't had to do for a VERY long time, restore a database off of an EMC Data Domain. Thankfully I wasn't restoring a failed production system, I was restoring to a replacement production system, so I was getting log shipping setup.

I've worked in plenty of shops with Data Domains before, but apparently I've blocked out the memories of doing a restore on them. Because if your backups are done the way EMC wants them to be done to get the most of the Data Domain (uncompressed SQL Backups in this case) the restore process is basically unusable. The reason that we are backing up the databases uncompressed was the allow the Data Domain to dedupe the backups as much as possible so that the final backup stored on the Data Domain would be as small as possible so that it could be replicated to another Data Domain in another data center.

The database in this case is ~6TB in size, so it's a big database. Running the restore off of the EMC Data Domain, was painfully slow. I canceled it after about 24 hours. It was at ~2% complete. Doing a little bit of math that database restore was going to take 25 days. While the restore was running we tried calling EMC support to see if there was a way to get the EMC Data Domain to allow the restores to run faster, and they answer was no, that's as fast as it'll run.

After stopping the restore, I backed up the same database to a local disk, and restored it to the new server from there. This time the restore took ~8 hours to complete. A much more acceptable number.

DCAC

If you are using EMC's Data Domain (or any backup appliance) do not use that appliance as your only location of your SQL Server backups. These appliances are very efficient at writing backups to them, and replicating those backups off to another site (which is what is being done in this case). But they are horrible at rehydrating those backups so that you can actually restore them. The proof of this is in the throughput of the restore commands. Here's the output of some of the restore commands that were running. These are for full backups, so there's nothing for SQL Server to process here, it's just moving blocks from point A to point B.

RESTORE DATABASE successfully processed 931 pages in 6.044 seconds (1.203 MB/sec).
RESTORE DATABASE successfully processed 510596 pages in 1841.175 seconds (2.166 MB/sec).
RESTORE DATABASE successfully processed 157903 pages in 440.849 seconds (2.798 MB/sec).
RESTORE DATABASE successfully processed 2107959 pages in 4696.428 seconds (3.506 MB/sec).
RESTORE DATABASE successfully processed 77307682 pages in 118807.557 seconds (5.083 MB/sec).
RESTORE DATABASE successfully processed 352411 pages in 816.810 seconds (3.370 MB/sec).
RESTORE DATABASE successfully processed 8400718 pages in 23940.799 seconds (2.741 MB/sec).
RESTORE DATABASE successfully processed 51554 pages in 111.890 seconds (3.599 MB/sec).
RESTORE DATABASE successfully processed 1222431 pages in 3167.605 seconds (3.014 MB/sec).

The biggest database there was restoring at 5 Megs a second. That was 33 hours to restore a database which is just ~606,816 Megs (~592 Gigs) in size. Now before you blame the SQL Server's or the network, all these servers are physical servers running on Cisco UCS hardware. The network is all 10 Gig networking, and the storage on these new servers is a Pure storage array. The proof that the network and storage was fine was the full restore of the database which was done from the backup to disk, as that was restored off of a UNC path which was still attached to the production server.

When testing these appliances, make sure that doing restores within an acceptable time window is part of your testing practice. If we had found this problem during a system down situation, the company would probably have just gone out of business. There's no way the business could have afforded to be down for ~25 days waiting for the database to restore.

Needless to say, as soon as this problem came up, we provisioned a huge LUN to the servers to start writing backups to. We'll figure out how to get the backups offsite (the primary reason that the Data Domain exists in this environment) another day (and in another blog post).

Denny

https://www.dcac.com/2015/07/20/if-you-thought-database-restores-were-slow-try-restoring-from-an-emc-data-domain/

백업 복구 인식 변화

단순한 데이터 보호에서 초고속 복구를 통한 비즈니스 보호로의 백업 개념 변화

재래식 백업 시스템

- 단순한 데이터 보호 목적
- 데이터 손실에 대한 보험
- 느린 백업 윈도우 및 복구 불확실성

현대적 데이터 보호

- **데이터 보호 → 비즈니스 보호**
- **데이터 활용 → 새로운 가치 창출**
- **초고속 성능 기반 복구 중심의 접근**

데이터가 아닌 비즈니스를 보호하려면…

1

사이버공격 대응

랜섬웨어 데이터 보호
신속한 업무 정상화
향상된 데이터 보안

2

빠른 복구 성능

비즈니스 보호
초고속 백업/복구
데이터 활용성 향상

3

서비스 무중단

AADC 구현
실시간 데이터 이중화
물리적 장애 보호

비즈니스 보호를 위한 새로운 패러다임

1

SafeMode

랜섬웨어 데이터 보호
신속한 업무 정상화
향상된 데이터 보안

2

RapidRestore

비즈니스 보호
초고속 백업/복구
데이터 활용성 향상

3

ActiveCluster

AADC 구현
실시간 데이터 이중화
물리적 장애 보호

사이버 공격 대응과 복구 중심 데이터 보호 전략

이제는 사이버 공격의 시대

피싱

휴먼 에러

자연 재해

랜섬웨어

CEO 사칭

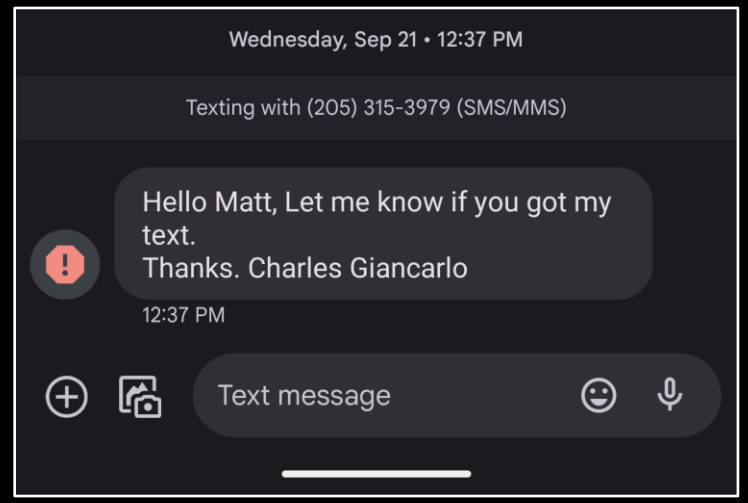

점점 커지는 사이버 공격의 위협

매년 더 많은 기업과 조직이 타겟이 되고 있습니다.

"88%가 시간당 4억원 이상의 비용 발생"

매년 증가하는 랜섬웨어 공격

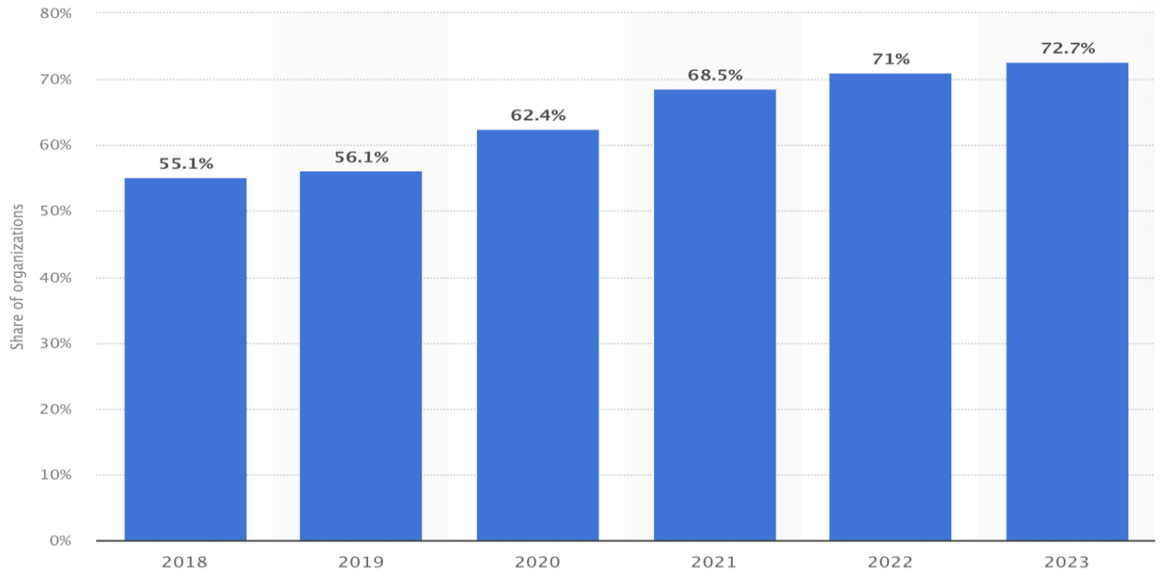

https://www.statista.com/statistics/204457/businesses-ransomware-attack-rate/

가동 중단에 따른 시간당 비용

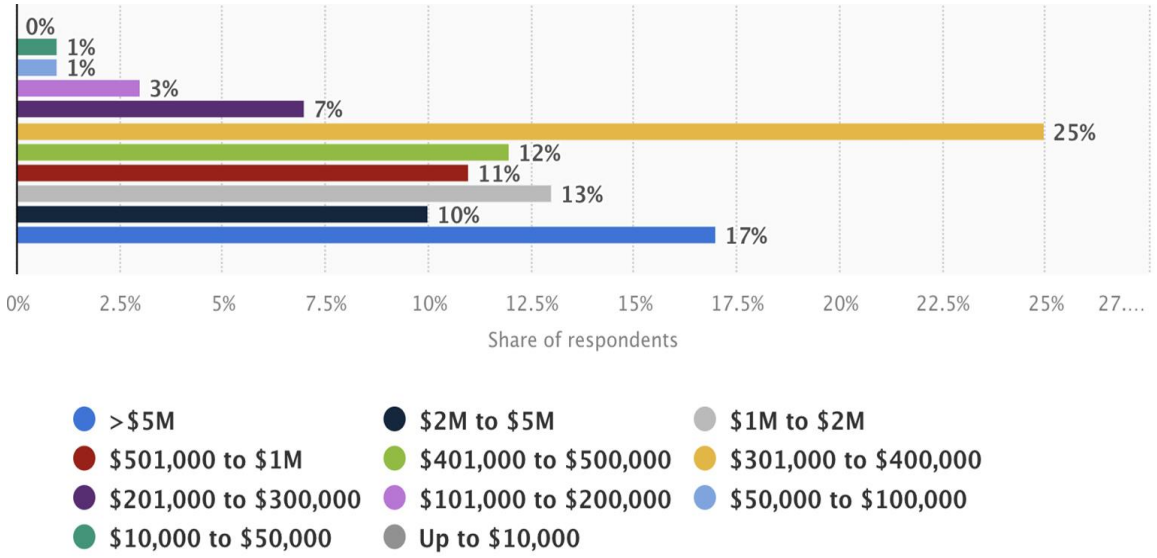

https://www.statista.com/statistics/753938/worldwide-enterprise-server-hourly-downtime-cost/

RANSOM

몸값, 몸값을 치르고 석방됨

MAL WARE

악성 소프트웨어

일정 금액을 지불 할 때까지 컴퓨터 시스템에 대한 액세스를 차단하도록 설계된 일종의 악성 소프트웨어

랜섬웨어 공격 구조

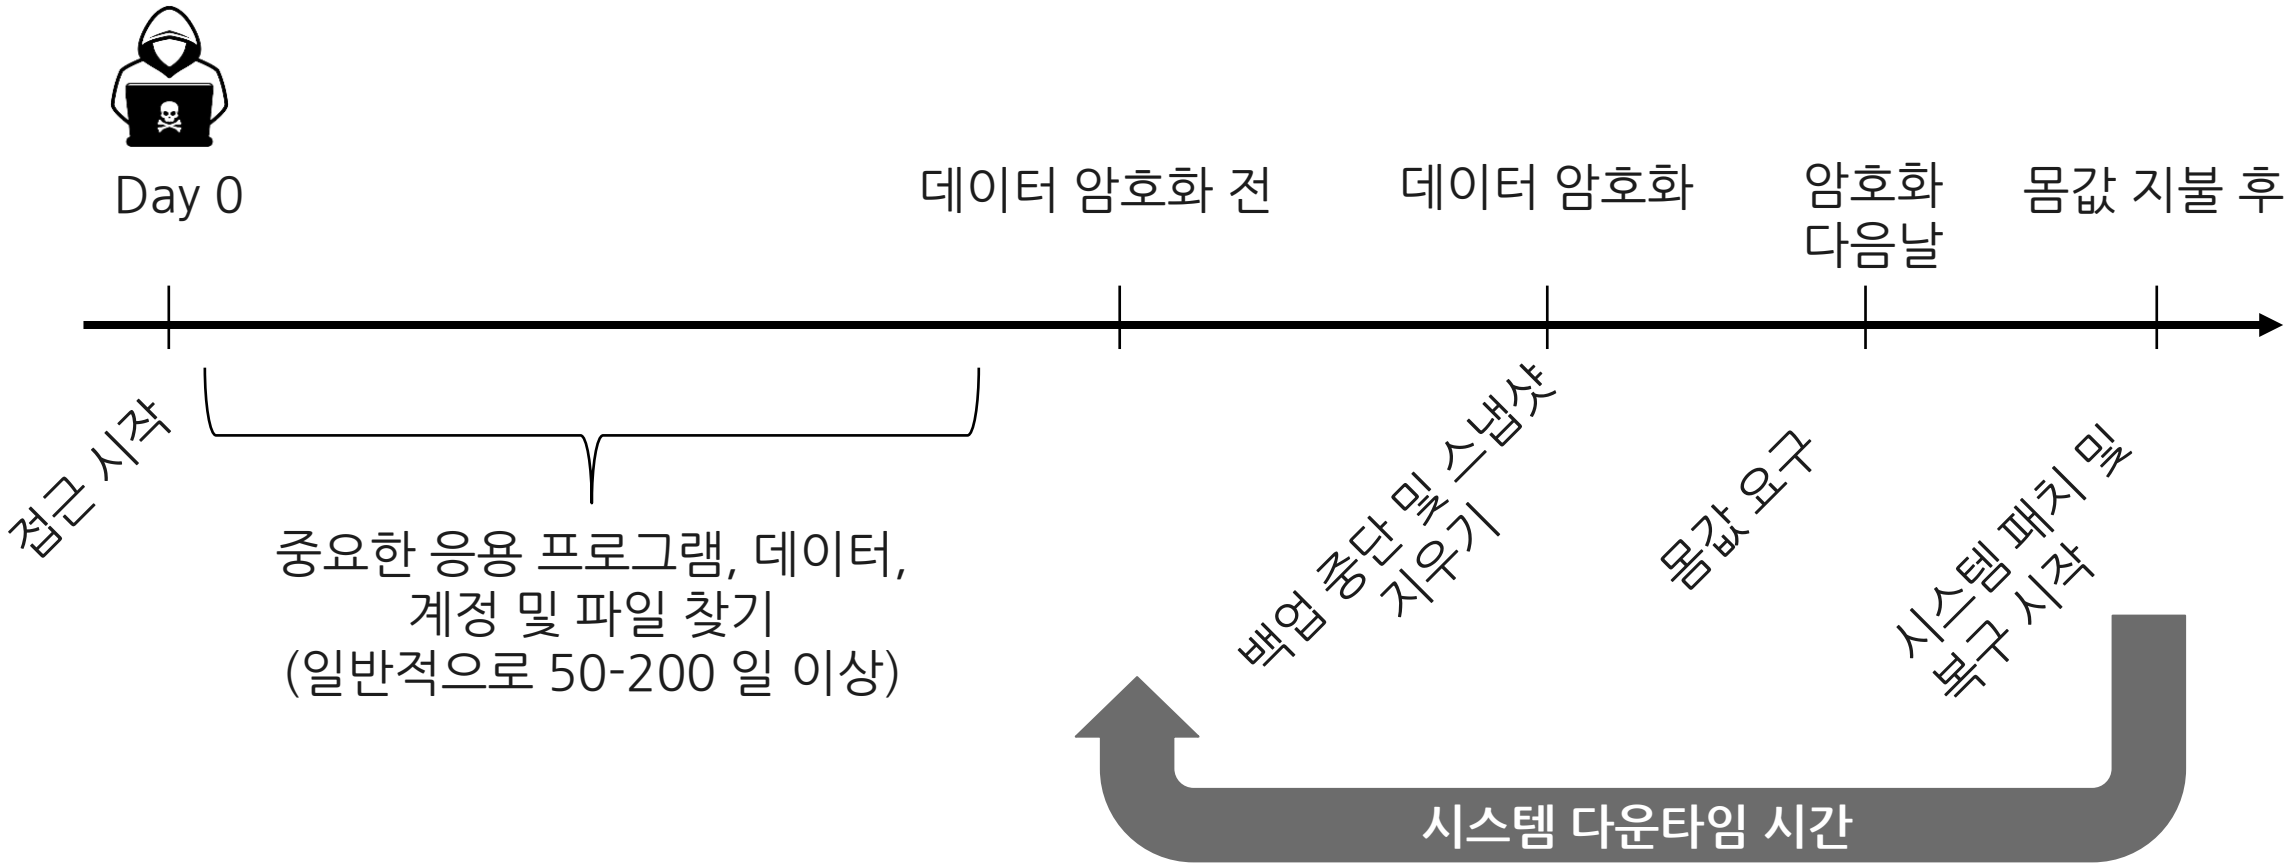
Day 0
데이터 암호화 전
데이터 암호화
암호화 다음날
몸값 지불 후
접근 시작
중요한 응용 프로그램, 데이터, 계정 및 파일 찾기 (일반적으로 50-200 일 이상)
백업 중단 및 스냅샷 지우기
몸값 요구
시스템 패치 및 복구 시작
시스템 다운타임 시간

해커는 항상 당신의 백업을 찾고 있습니다

97% 백업본 감염을 위한 랜섬웨어 공격 시도

73% 백업본에 대한 랜섬웨어 공격 성공

36% 몸값을 지불했으나 데이터 복구 실패

Source: Veeam 2022년 랜섬웨어 동향 보고서

공격을 받았다면 다음 두 가지 대응이 필요합니다.

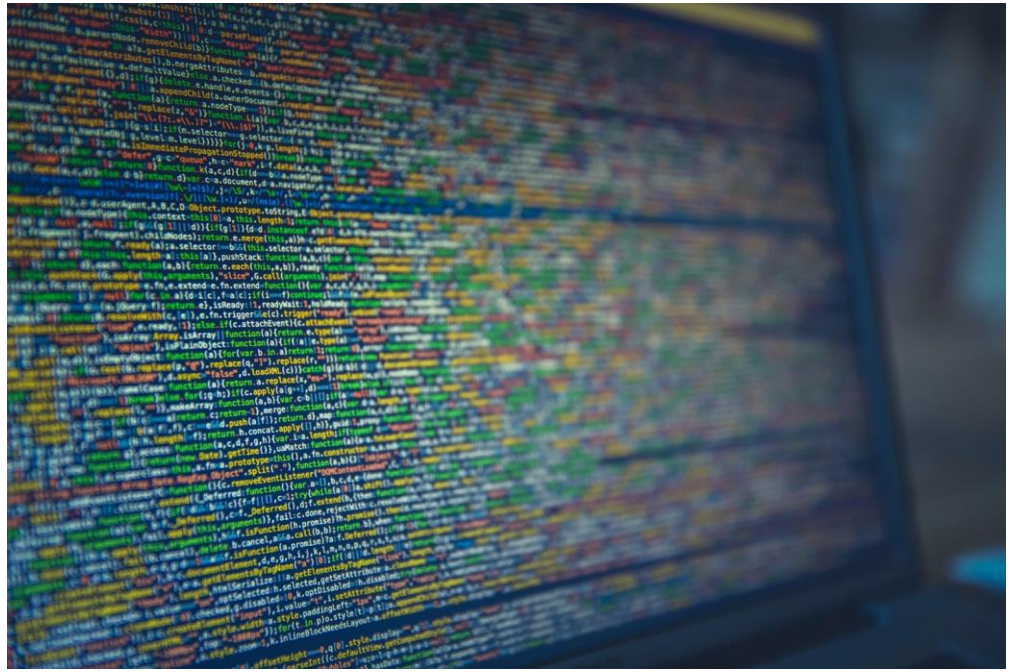

랜섬웨어 공격에도
유효하고 사용가능한 데이터 복사본

대량의 데이터에 대한
초고속 데이터 복구

#1. 랜섬웨어 공격에도 유효하고 사용가능한 데이터 복사본

Snapshot Policy	Authorization	Tune Eradication Timer	Disable Eradication
위/변조 불가능한 스냅샷	권한 있는 사용자 제한	완전 삭제 타이머 설정	변경되지 않는 안전한 데이터
유연하고 세분화된 스냅샷 정책	최대 5명까지 승인된 컨택포인트, PIN code 제공	24시간에서 최대 30일까지 스냅샷 보관	볼륨 수동 완전삭제 비활성화

백업 소프트웨어와의 협업

- 고성능 백업 타겟 스토리지
- VM 단위 백업 성능 개선
- 백업 스트림 병렬 수행
- 백업본 에 대한 SafeMode 적용
  → 백업 위변조 및 삭제 방지

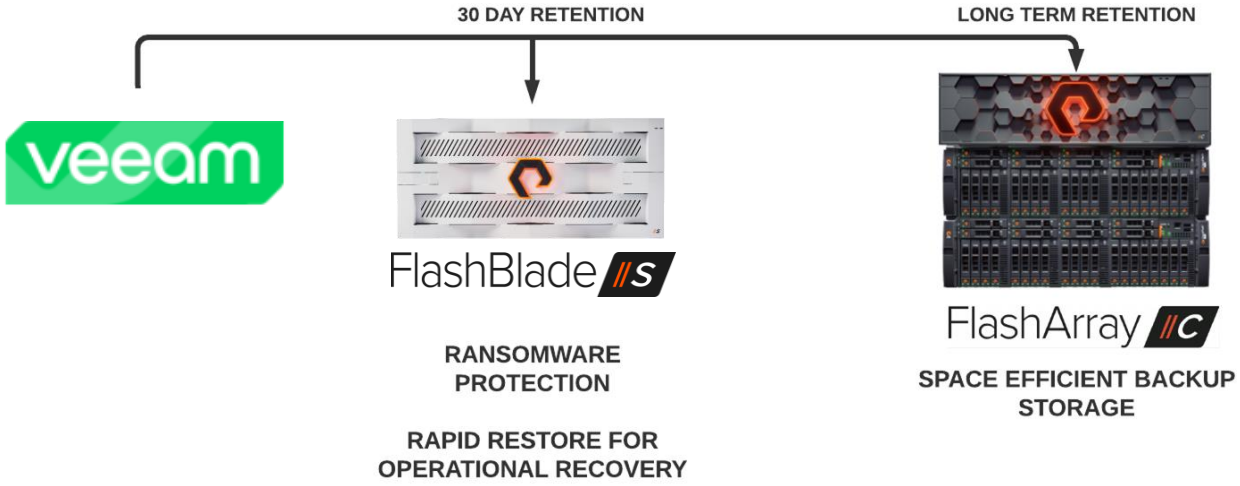

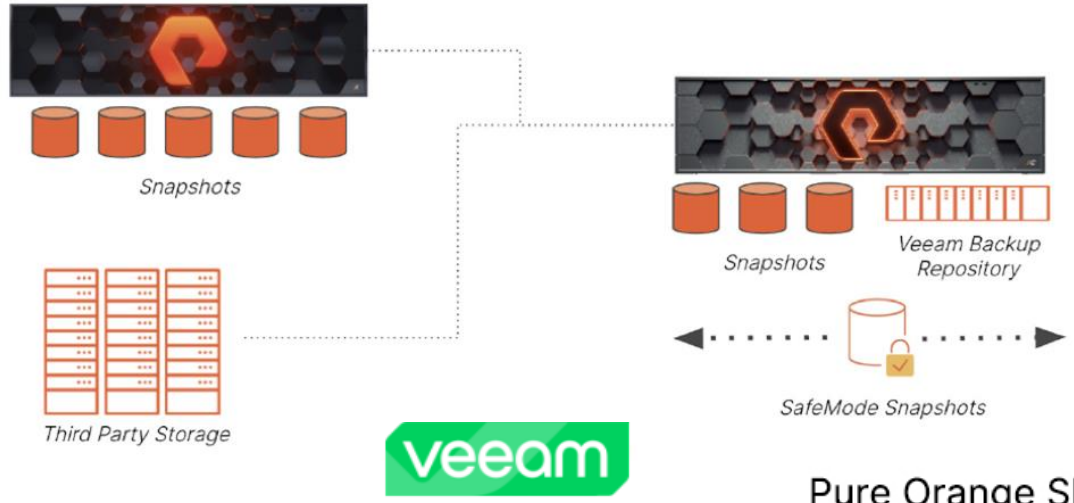

DRR Anomaly Detection

PURE1 을 통해 데이터 위변조 시,
데이터 절감율 변화 추이를 분석하여
사이버공격에 대한 이벤트 발생

이를 통해 운영자가 신속하게 재해
상황을 대응할 수 있도록 지원

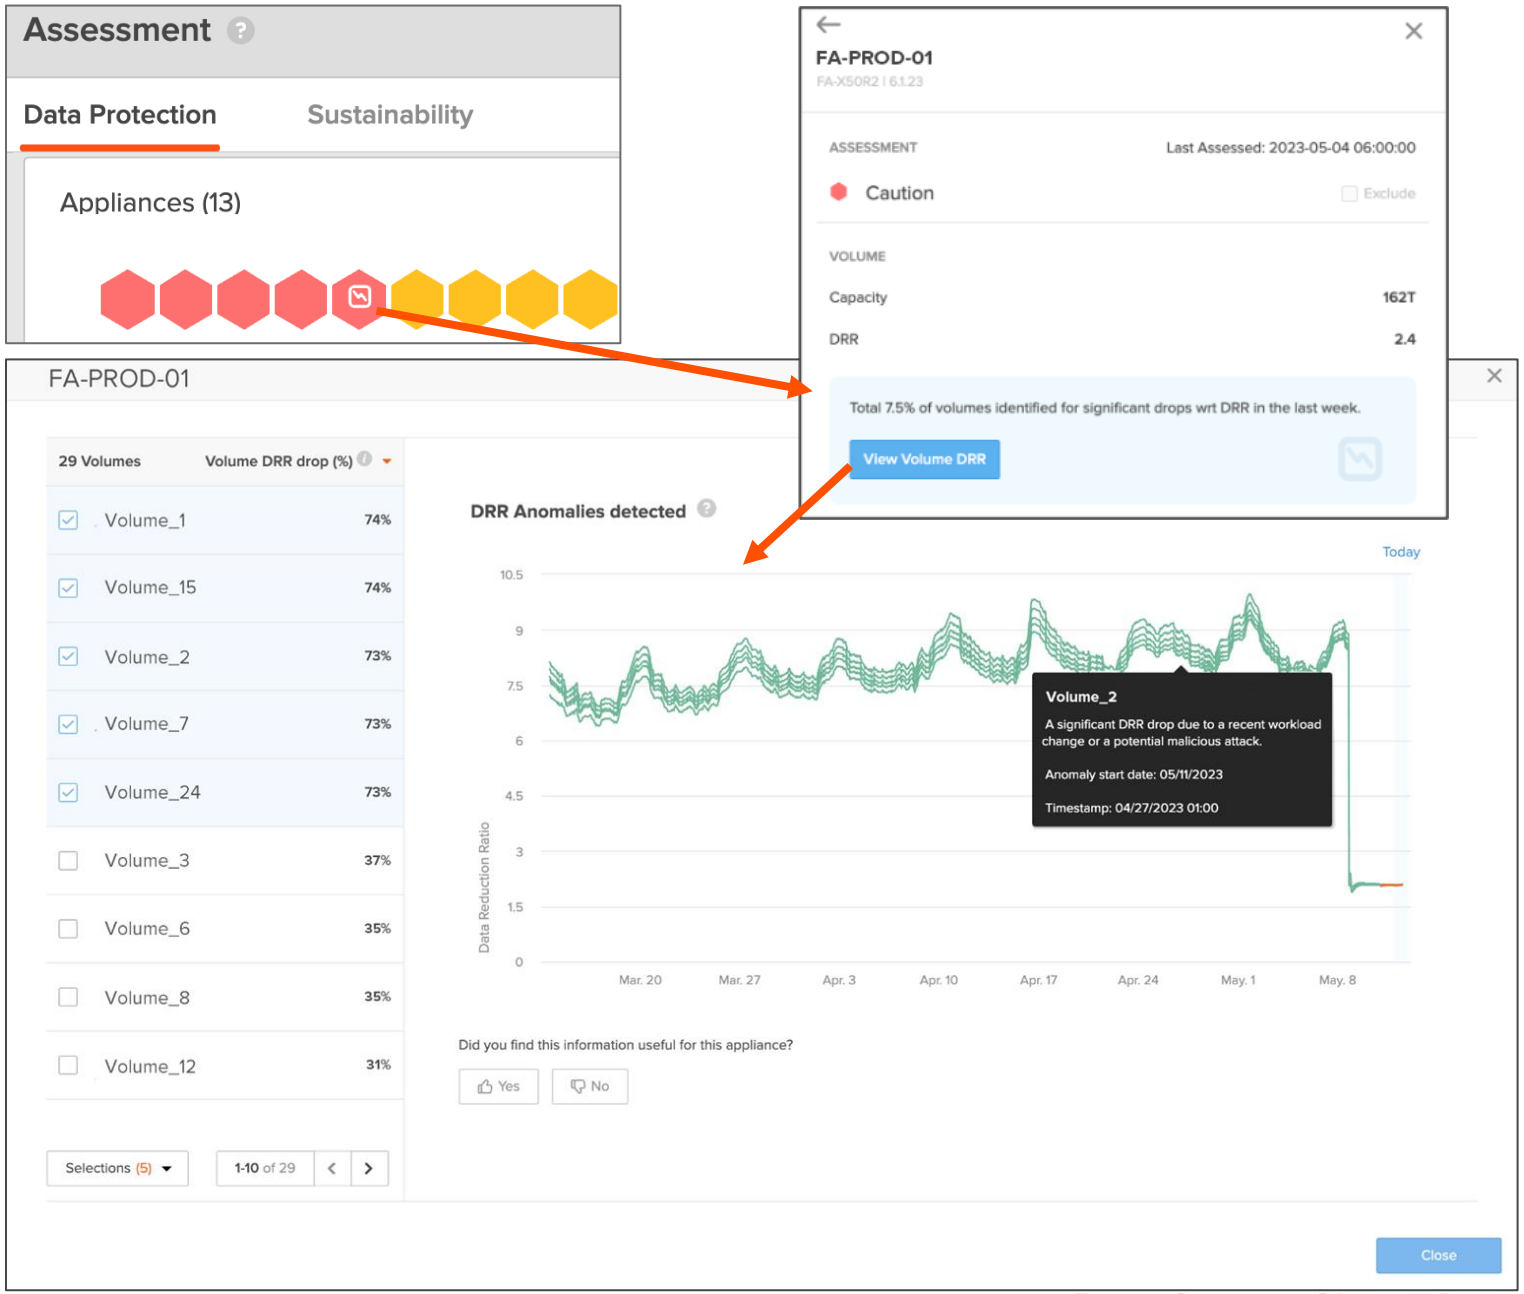

사이버 공격 대응 실제 고객 사례

- 실제 해킹으로 인해 스토리지 내 스냅샷 삭제 오퍼레이션 수행

Status : Ransomware attack occured on Jun 14th.

2022-06-14 02:59:54 | 1949506 | customer | pureuser | purepgroup destroy
2022-06-14 03:00:00 | 1949507 | customer | pureuser | purepgroup destroy
2022-06-14 03:00:06 | 1949512 | customer | pureuser | purepgroup destroy

스냅샷 삭제

Eradication Bucket
- 삭제된 볼륨/스냅샷 임시 저장 영역
- 정해진 기간 내 임의 삭제 불가 (관리자 포함)
- 해커 공격 시, 즉시 활용하여 데이터 시점 복구

해커에 의해
데이터 암호화 수행

3일 후 업무 정상화 → **SafeMode 기간 5일로 연장**

#2. 대량의 데이터에 대한 초고속 데이터 복구

재래식 백업은 왜 복구가 느릴까?

최대 처리량	최대 14.0TB/시간
논리적 용량[1,2]	최대 14.1PB
Cloud Tier 사용 시	최대 42.2PB
가용 용량	12TB~256TB
Cloud Tier 사용 시	최대 768TB

처리량 ≠ 복구성능

https://www.delltechnologies.com/asset/ko-kr/products/data-protection/briefs-summaries/powerprotect-data-manager-appliance-ds.pdf

재래식 백업은 비즈니스를 보호할 수 있나요?

DCAC

If You Thought Database Restores Were Slow, Try Restoring From an ~~EMC Data Domain~~

By: Denny Cherry

Published On: July 20, 2015

Recently I did something which I haven't had to do for a VERY long time, restore a database off of an EMC Data Domain. Thankfully I wasn't restoring a failed production system, I was restoring to a replacement production system, so I was getting log shipping setup.

I've worked in plenty of shops with Data Domains before, but apparently I've blocked out the memories of doing a restore on them. Because if your backups are done the way EMC wants them to be done to get the most of the Data Domain (uncompressed SQL Backups in this case) the restore process is basically unusable. The reason that we are backing up the databases uncompressed was the allow the Data Domain to dedupe the backups as much as possible so that the final backup stored on the Data Domain would be as small as possible so that it could be replicated to another Data Domain in another data center.

The database in this case is ~6TB in size, so it's a big database. Running the restore off of the EMC Data Domain, was painfully slow. I canceled it after about 24 hours. It was at ~2% complete. Doing a little bit of math that database restore was going to take 25 days. While the restore was running we tried calling EMC support to see if there was a way to get the EMC Data Domain to allow the restores to run faster, and they answer was no, that's as fast as it'll run.

After stopping the restore, I backed up the same database to a local disk, and restored it to the new server from there. This time the restore took ~8 hours to complete. A much more acceptable number.

DCAC

If you are using EMC's Data Domain (or any backup appliance) do not use that appliance as your only location of your SQL Server backups. These appliances are very efficient at writing backups to them, and replicating those backups off to another site (which is what is being done in this case). But they are horrible at rehydrating those backups so that you can actually restore them. The proof of this is in the throughput of the restore commands. Here's the output of some of the restore commands that were running. These are for full backups, so there's nothing for SQL Server to process here, it's just moving blocks from point A to point B.

RESTORE DATABASE successfully processed 931 pages in 6.044 seconds (1.203 MB/sec).
RESTORE DATABASE successfully processed 510596 pages in 1841.175 seconds (2.166 MB/sec).
RESTORE DATABASE successfully processed 157903 pages in 440.849 seconds (2.798 MB/sec).
RESTORE DATABASE successfully processed 2107959 pages in 4696.428 seconds (3.506 MB/sec).
RESTORE DATABASE successfully processed 77307682 pages in 118807.557 seconds (5.083 MB/sec).
RESTORE DATABASE successfully processed 352411 pages in 816.810 seconds (3.370 MB/sec).
RESTORE DATABASE successfully processed 8400718 pages in 23940.799 seconds (2.741 MB/sec).
RESTORE DATABASE successfully processed 51554 pages in 111.890 seconds (3.599 MB/sec).
RESTORE DATABASE successfully processed 1222431 pages in 3167.605 seconds (3.014 MB/sec).

The biggest database there was restoring at 5 Megs a second. That was 33 hours to restore a database which is just ~606,816 Megs (~592 Gigs) in size. Now before you blame the SQL Server's or the network, all these servers are physical servers running on Cisco UCS hardware. The network is all 10 Gig networking, and the storage on these new servers is a Pure storage array. The proof that the network and storage was fine was the full restore of the database which was done from the backup to disk, as that was restored off of a UNC path which was still attached to the production server.

When testing these appliances, make sure that doing restores within an acceptable time window is part of your testing practice. If we had found this problem during a system down situation, the company would probably have just gone out of business. There's no way the business could have afforded to be down for ~25 days waiting for the database to restore.

Needless to say, as soon as this problem came up, we provisioned a huge LUN to the servers to start writing backups to. We'll figure out how to get the backups offsite (the primary reason that the Data Domain exists in this environment) another day (and in another blog post).

Denny

https://www.dcac.com/2015/07/20/if-you-thought-database-restores-were-slow-try-restoring-from-an-emc-data-domain/

퓨어스토리지 비즈니스 보호 플랫폼 “백업 복구” 성능

쓰기(백업) / 읽기(복구) 성능

Solution	FlashArray //C	FlashArray //E	FlashBlade //S200	FlashBlade //E
Commvault	SMB (~10-15 TB/hr)	SMB (2.7TB/hr write, 5.4 TB/hr read max)	Object (35-50 TB/hr) read (15-30 TB/hr) write	Object (30 TB/hr) read (10 TB/hr) write
veeam	Block (~10-15 TB/hr)	SMB (2.7TB/hr write, 5.4 TB/hr read max)	Object (35-50 TB/hr) read (15-30 TB/hr write)	Object (30 TB/hr) read (10 TB/hr write)
VERITAS	Block (~10-15 TB/hr)	SMB (2.7TB/hr write, 5.4 TB/hr read max)	NFS (35-50 TB/hr) read (15-30 TB/hr) write	Object (30 TB/hr) read (10 TB/hr write)
SQL ORACLE	TBA	SMB (2.7TB/hr write, 5.4 TB/hr read max)	SQL: SMB (47.9 TB/hr write, 63.3 TB/hr read) S3 (44.6 TB/hr write, 62.3 TB/hr read)	SQL: SMB: (21.8 TB/hr write, 64.3 TB/hr read) S3: (44.6 TB/hr write, 62.3 TB/hr read)

개별 백업 환경에 따라 성능이 상이할 수 있음

FB//E vs DD?

10x 빠른 복구 성능 **5.5x** 더 높은 용량 집적도

78% 상면 절감 효과 **72%** 에너지 소비 개선

- FlashBlade//E 3PB 용량 구성 예시

$1.1M

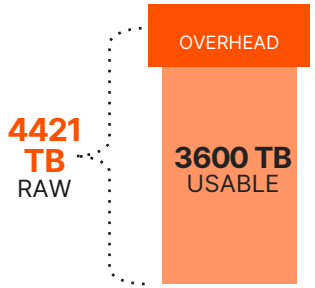

$4M-$5M

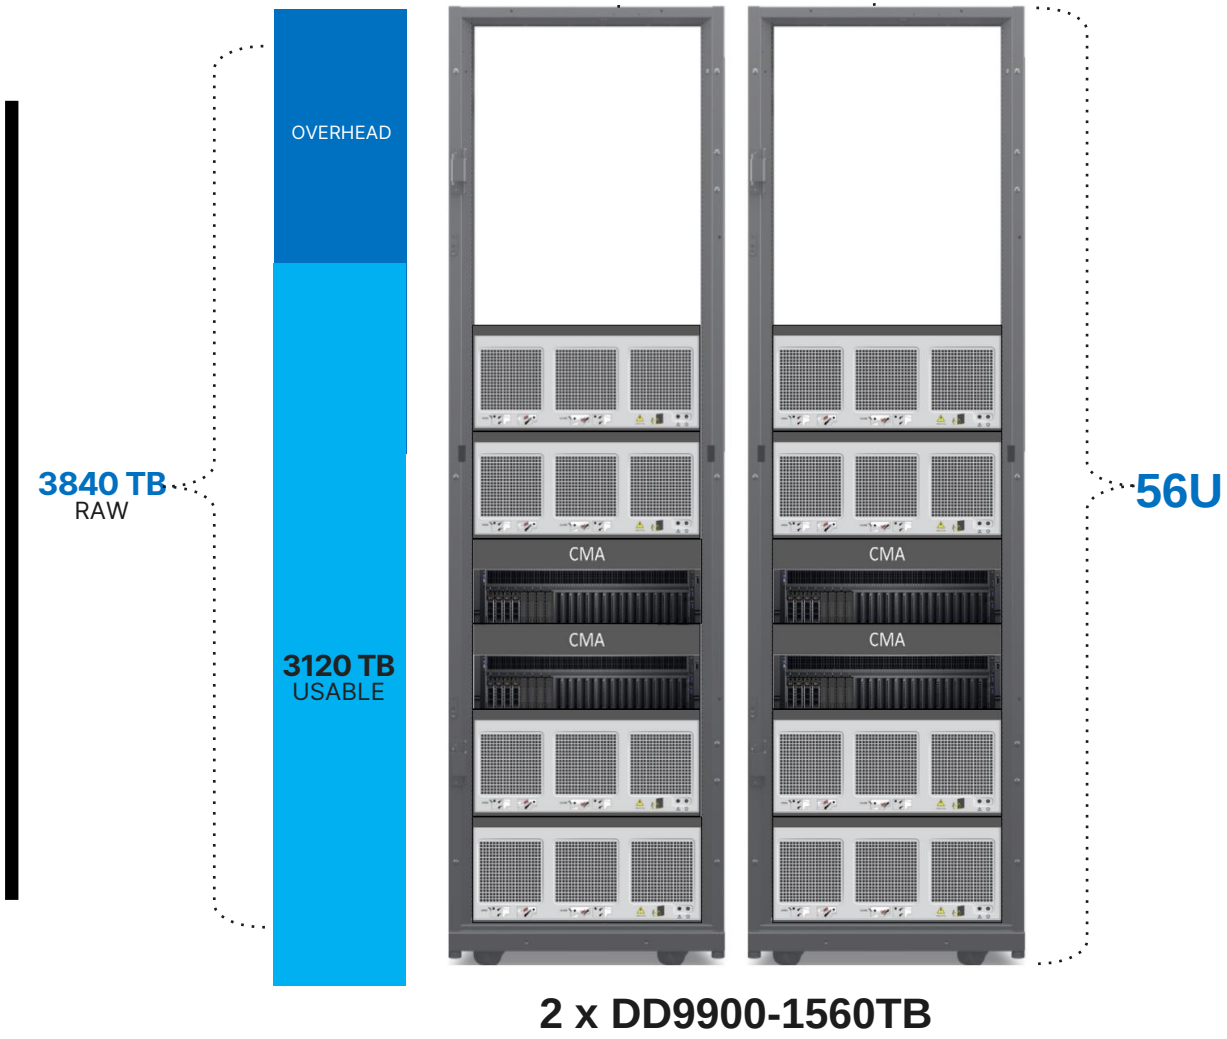

2 x DD9900-1560TB

적용 사례

1. 국내 대형 금융 플랫폼 기업 - VTL Replication 환경을 FlashArray//C 로 전환

   → 65% 상면 절감 / 70% TCO 절감

2. 인도 대형 금융 기업 - 기존 백업 VTL 환경을 FlashBlade//E 로 전환

   → 기존 대비 백업/복구 성능 10배 이상 개선

   → 복구 성능: 11+TB/Hr, 백업 성능: 44TB/Hr (5배 성능 개선)

   → 2개 랙 상면을 12RU 으로 통합

   → TCO 80% 개선

서비스 무중단

재해 복구 RPO/RTO 수준 정의

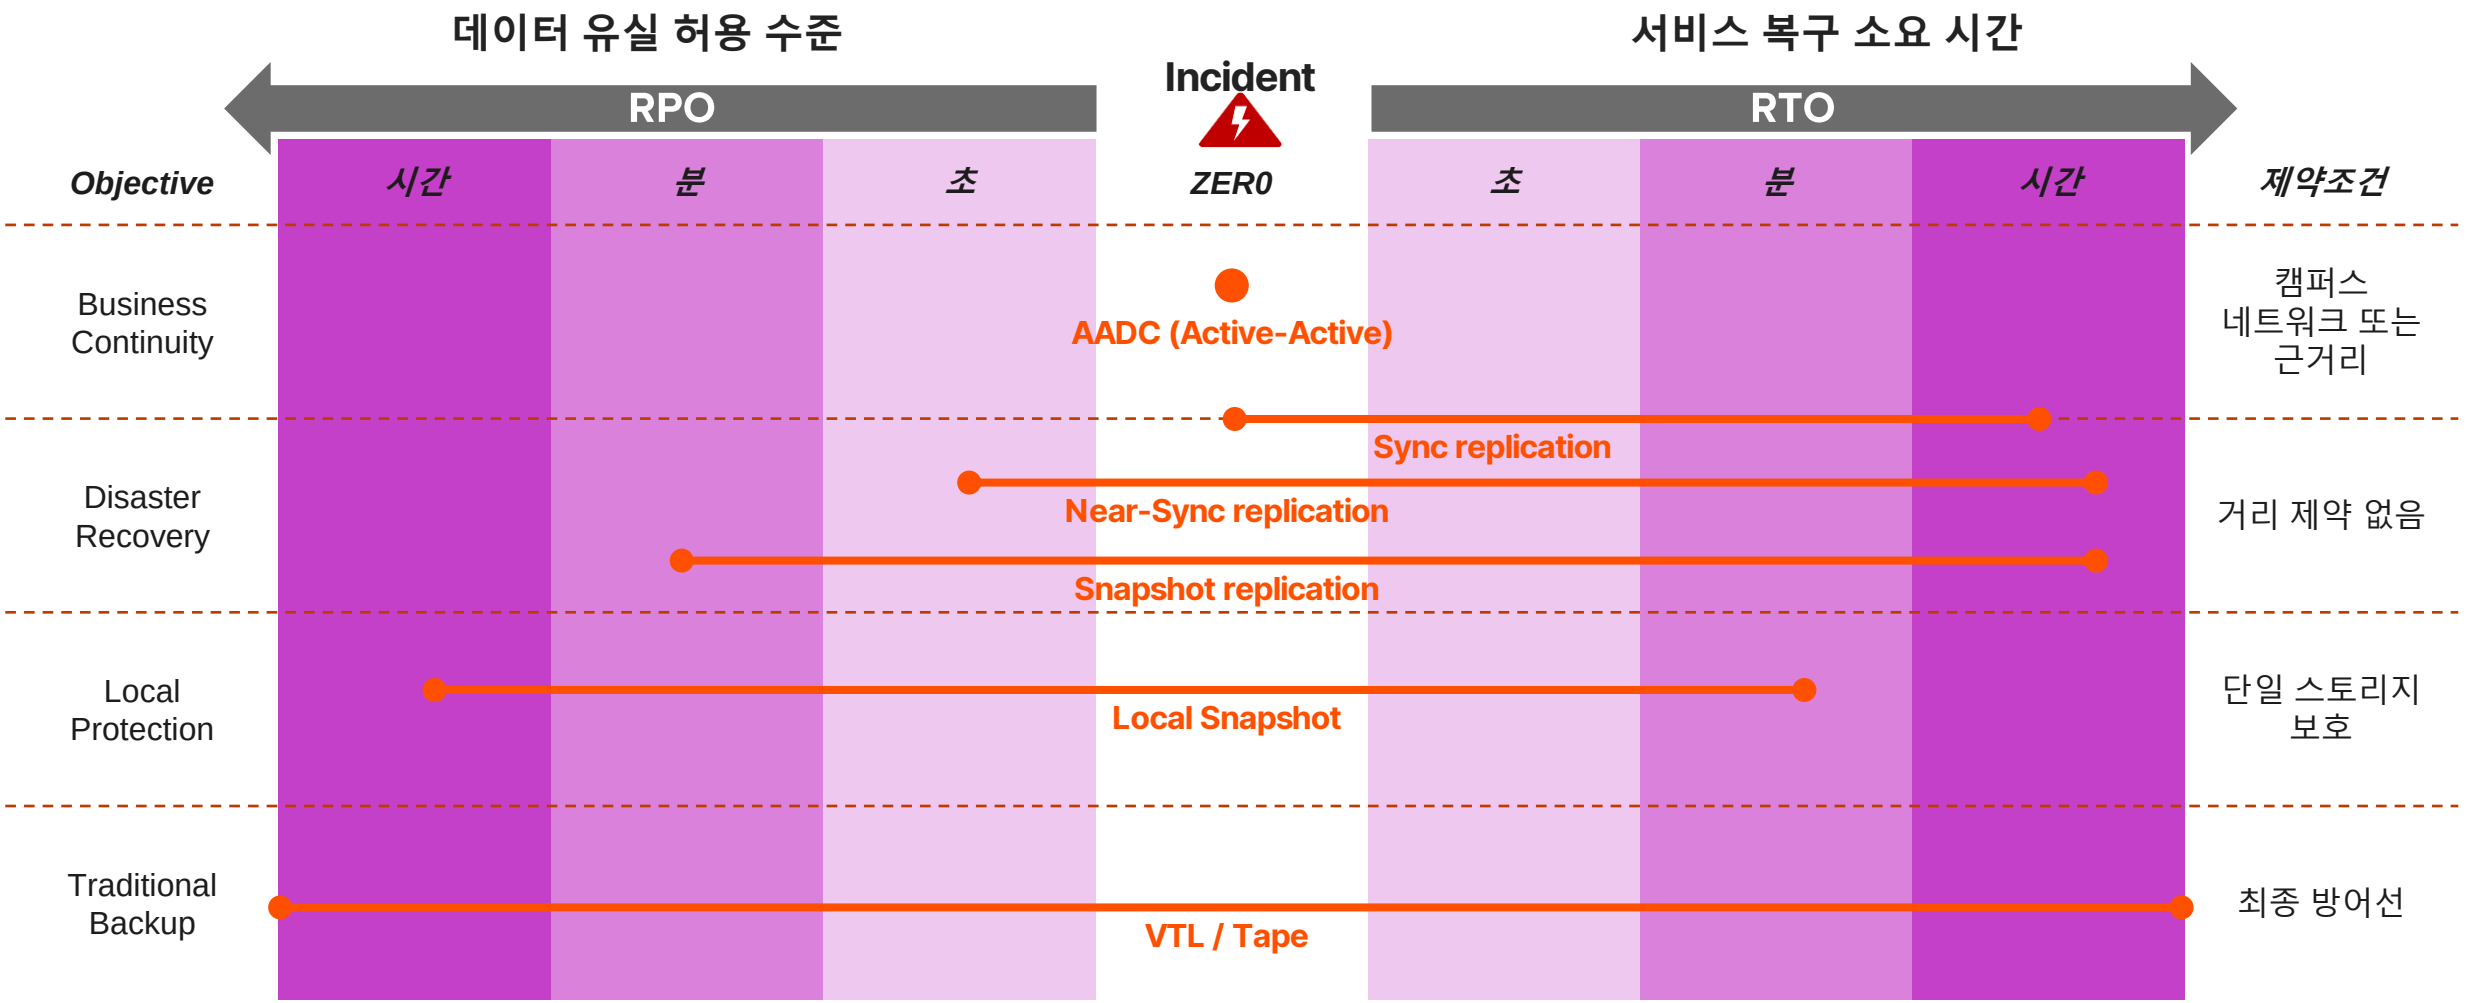

비즈니스 연속성을 위한 새로운 아키텍처

Active/Standby → Active/Active 로의 전환

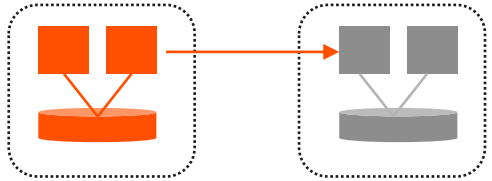

Active/Standby

복잡한 페일오버 절차

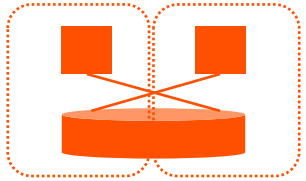

Active/Active

무중단 서비스 지원

DR 구성을 위한 스토리지 이중화 아키텍처 전환

ActiveCluster Overview

- **Active-Active 기반 데이터센터 장애 보호 RPO = Zero**
- **데이터 손실 없는 즉시 Failover RTO = Zero**
- **사용자 개입 없는 자동화된 장애 조치**
- **Round Trip Time(RTT) = 11ms 내에서의 구성**
- **Async 스냅샷 방식에서 Sync 모드로 자동 전환**
- **클라우드 기반 Mediator 제공으로 별도 Witness 불필요**

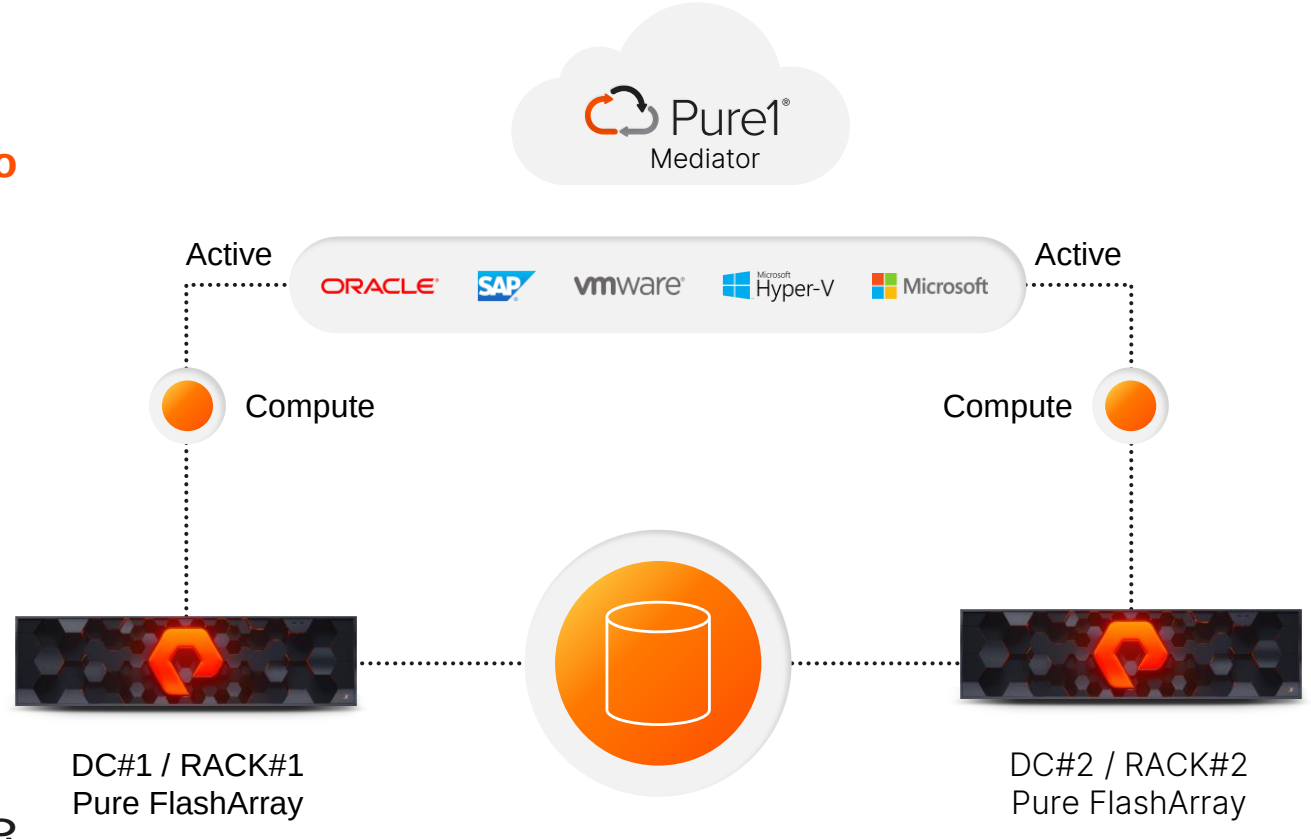

Multi-Site Active / Active

Zero RPO, Zero RTO, Zero $€¥£, Zero Additional HW

서비스 무중단 - AADC 구현 예시

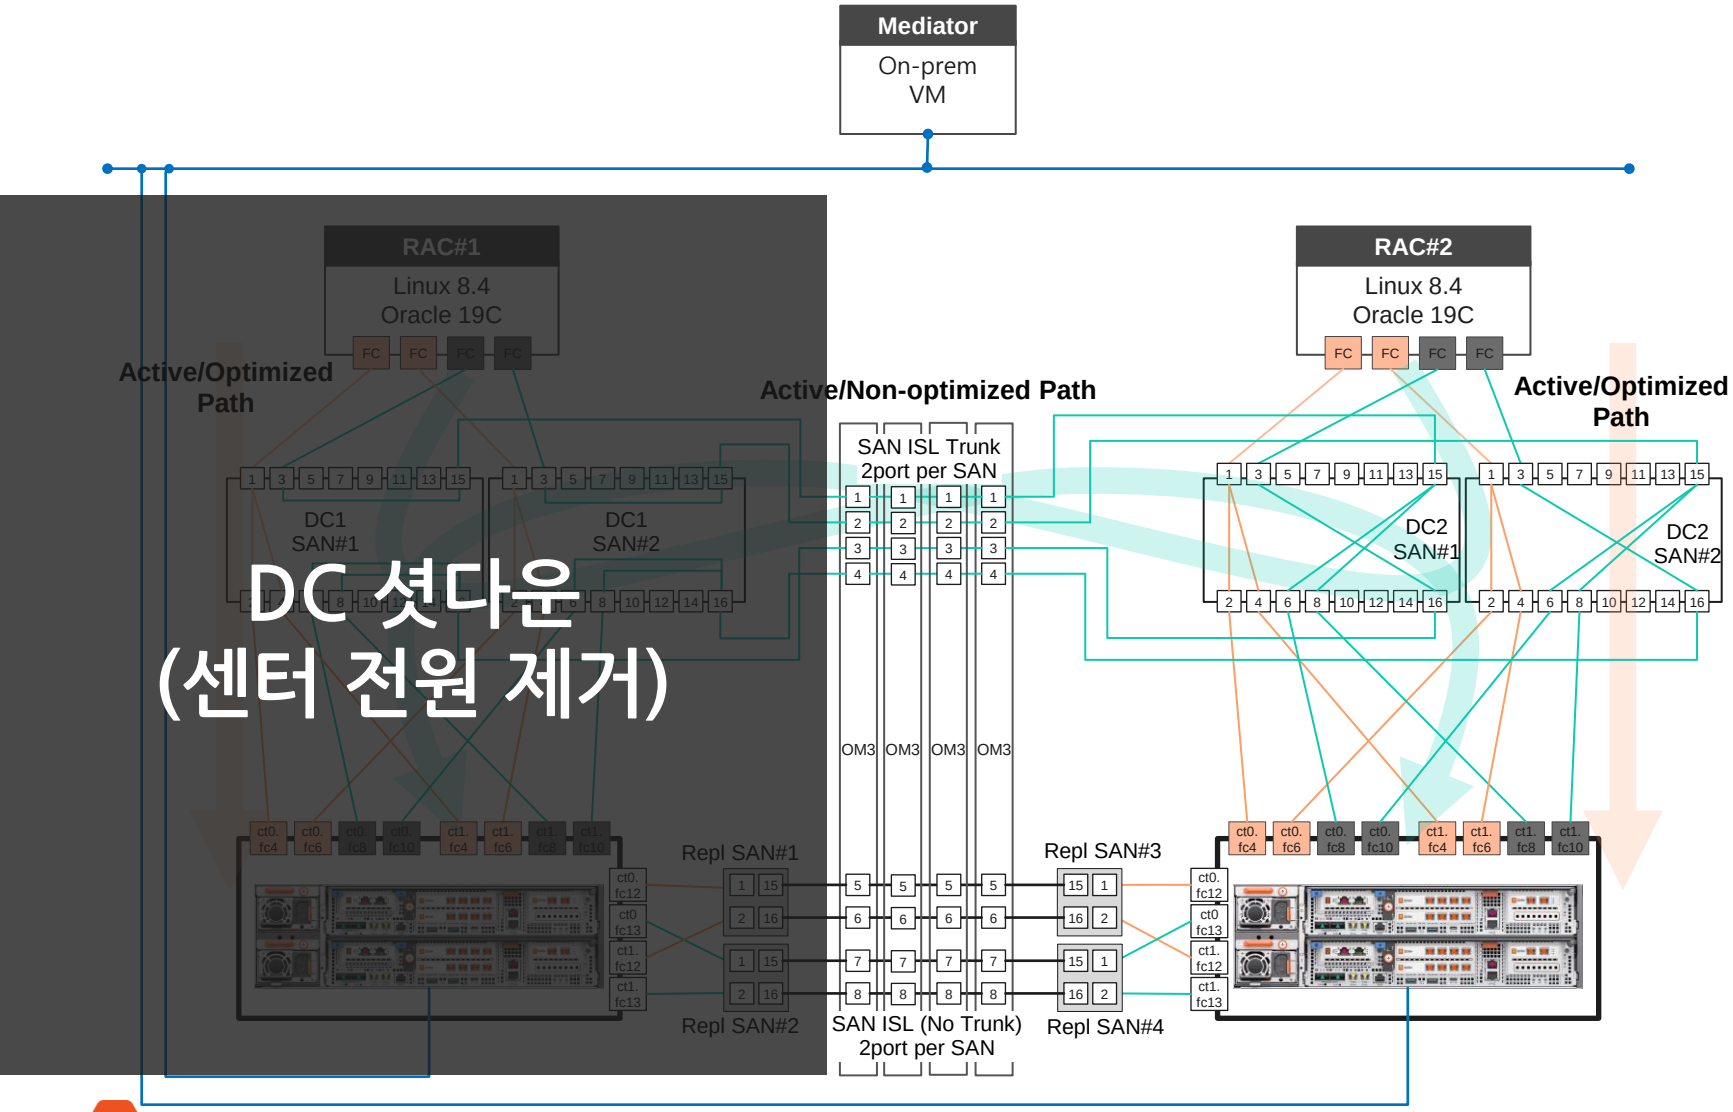

상세 내용

테스트 요약
- 센터간 Active-Active DC 구현
- Oracle Stretched RAC 구성
- 센터 장애 시, 무중단으로 서비스 전환 검증

검증 시나리오
- Write-intensive
- Oracle Swingbench 사용
- Update 위주의 데이터베이스 워크로드 검증
- 스토리지 전체 장애 / 센터 전제 장애 검증

대상 장비
- FlashArray//X20R4 x 2 SET

서비스 무중단 - AADC 구현 예시

전체 센터 장애 상황 구현 및 결과: 별도 운영자 개입없이 모든 서비스 및 사용자 세션 모두 정상적으로 무중단 전환

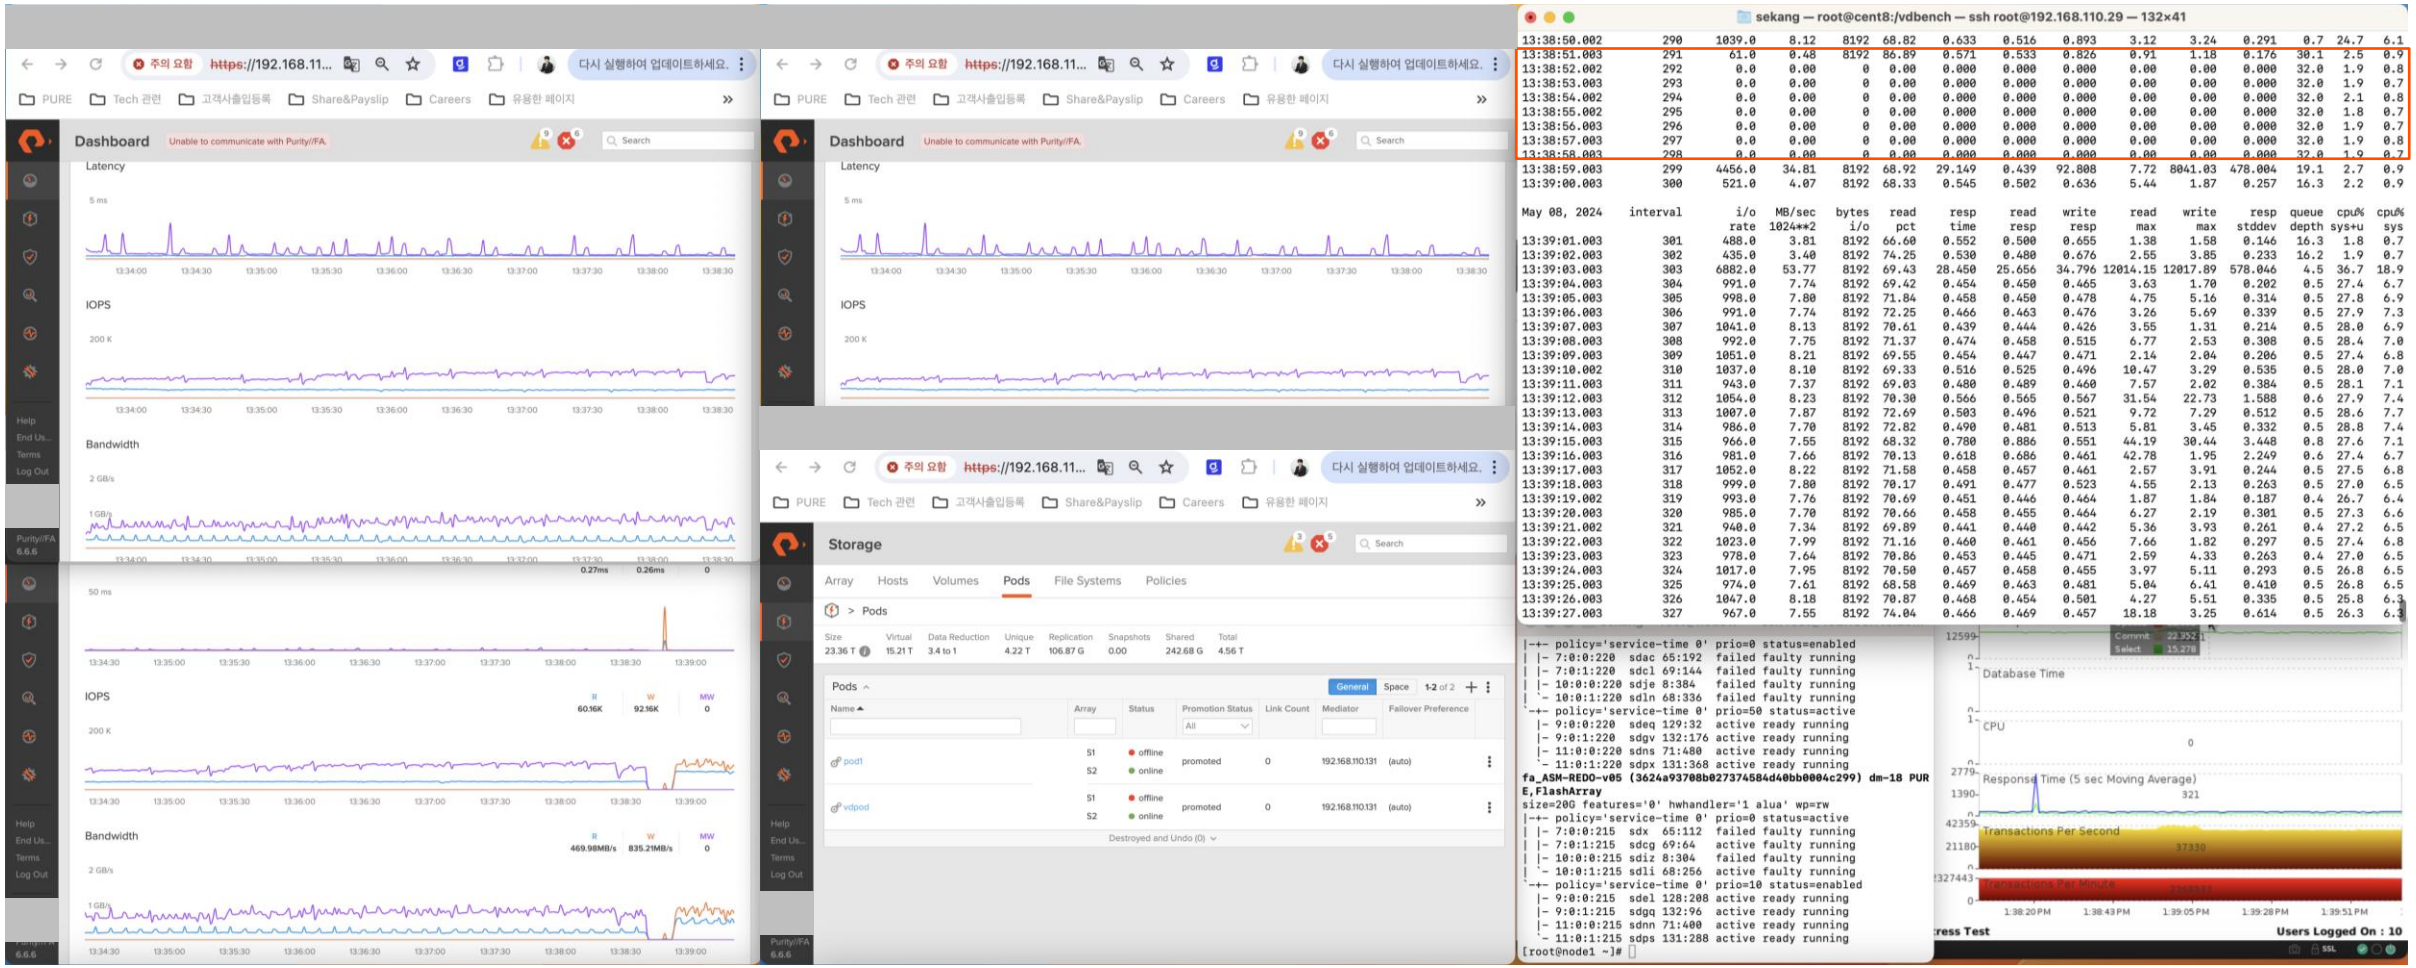

Executive Summary

사이버 공격 복구

SafeMode 를 통해
사이버 공격으로부터의
신속하게 복구!

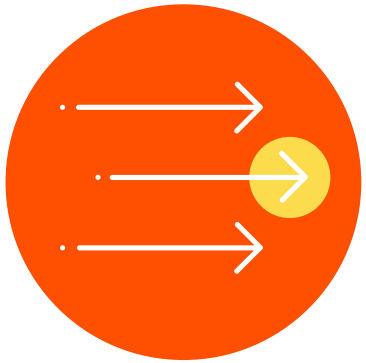

초고속 복구 성능

처리량이 아닌
진정한 복구 중심 성능으로
데이터 복구를 빠르게!

서비스 무중단

Active-Active DR 설계로
센터 장애 시에도
무중단으로 서비스 보호!

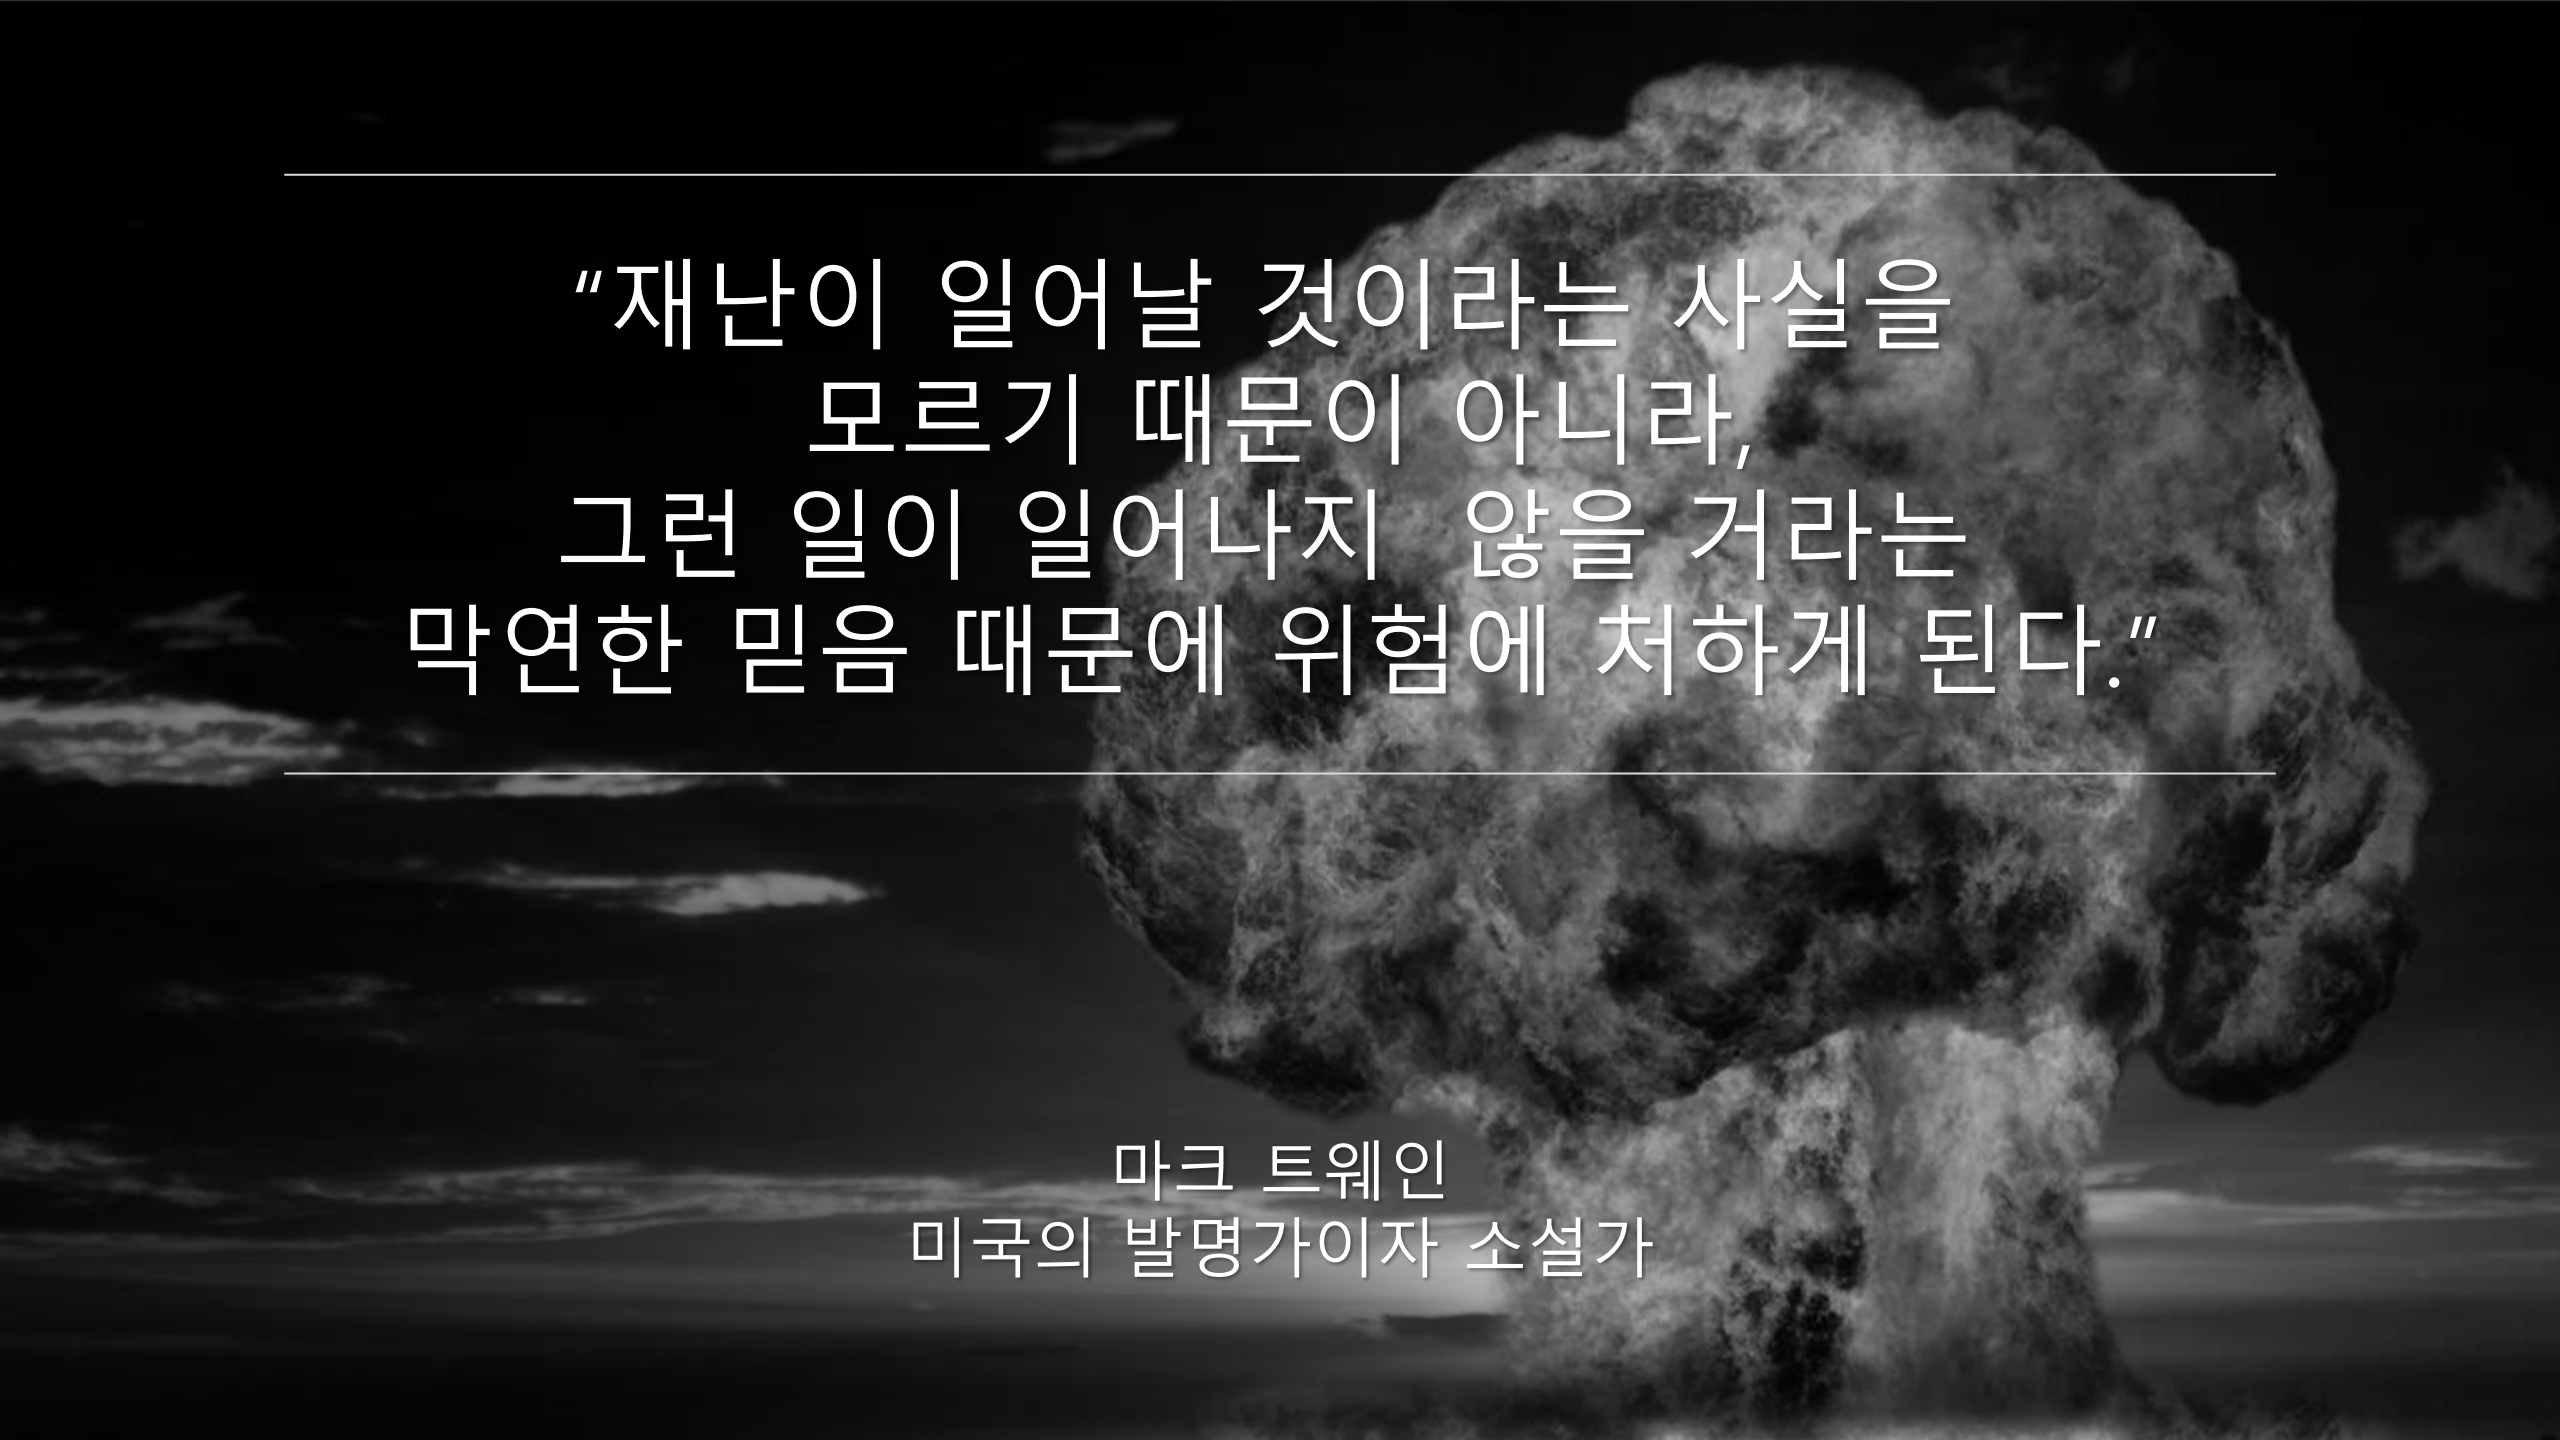
"재난이 일어날 것이라는 사실을
모르기 때문이 아니라,
그런 일이 일어나지 않을 거라는
막연한 믿음 때문에 위험에 처하게 된다."
마크 트웨인
미국의 발명가이자 소설가

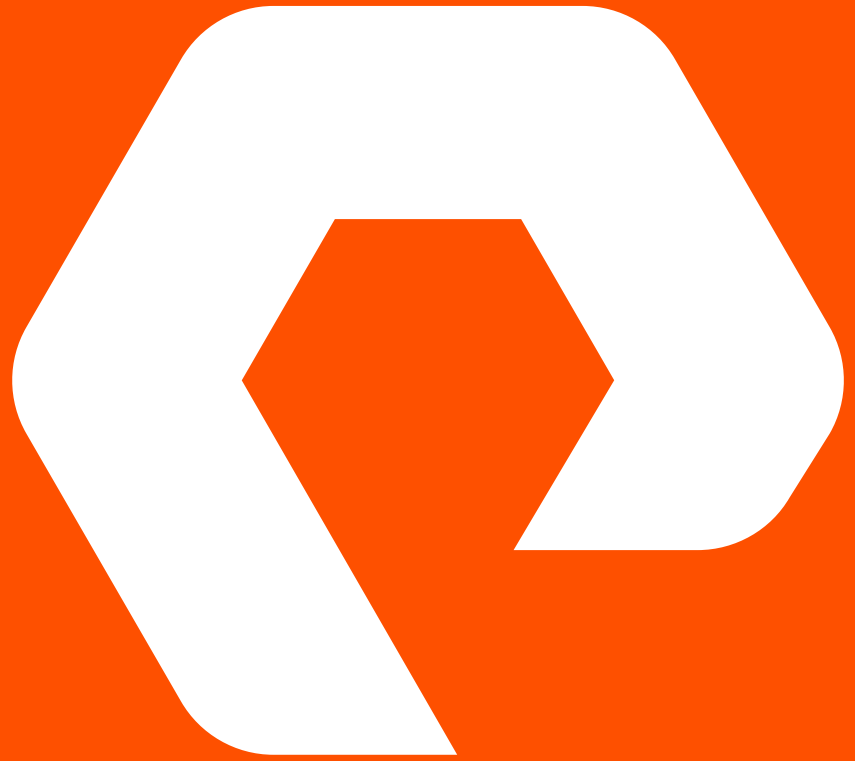

Uncomplicate Data Storage, Forever