WHITE PAPER

BROADCOM SOFTWARE

Migrating to Next-Generation DX APM

Key Options, Considerations, and Best Practices

Henrik Nissen Ravn
Principal Engineering Solution Architect and APM Champion
Broadcom Software

Table of Contents

Executive Summary

When moving to the newest version of Broadcom's DX Application Performance Management (DX APM) solution, customers have the options of deploying DX APM on-premises for a private cloud installation, or migrating to DX APM SaaS, which is DX APM offered as a public cloud service. The software base is identical for both deployment options. DX APM and DX APM SaaS are often collectively referred to as Broadcom's next-generation DX APM for this reason.

DX APM features a modern and agile micro-services architecture based on open source. The solution uses a unified data lake to provide normalized and correlated monitoring data for deep analytics and machine learning. This results in dramatically improved insights and enables more proactive, and potentially automated, remediation.

This architecture represents a significant advancement, offering a number of benefits to DX APM customers:

- Monitoring scalability, addressing enterprise needs for growth and full observability.
- First-class security capabilities that support sustained compliance and protection.
- Improved manageability, requiring fewer resources for upgrades, maintenance, and management.
- Increased flexibility in deploying Broadcom software and integrated third-party domain tools.
- Faster evolution of the DX platform and DX APM to meet changing customer demands and expectations.
- Increased efficiency in data processing, providing optimized resource utilization, response times, and capacity, and improved predictability and stability.

The solution uses a unified data lake to provide normalized and correlated monitoring data for deep analytics and machine learning. This results in dramatically improved insights and enables more proactive, and potentially automated, remediation.

Collectively, these architectural advancements allow Broadcom's next-generation DX APM, to be a secure, agile, and cost-effective SaaS or on-premises solution. DX APM offers these benefits regardless of on-premises deployment of DX APM or the public cloud offering of DX APM SaaS. The DX APM SaaS option eases the transition to DX APM by eliminating deployment, alleviating management, and reducing requirements for skills and resources.

This white paper offers a detailed look at the key considerations and options for making the move to DX APM. In addition, it offers a number of best practices for managing a migration successfully.

DX APM Microservices Architecture and Capacity

Key Architectural Changes

With the release of DX OI, the architecture of Broadcom's APM solution has changed substantially in order to enable deployment on the foundational DX platform. The new architecture is one based on clustered microservices, while APM is based on the monolithic architecture of the Enterprise Manager (EM).

The capabilities of the EM have been decomposed into small, independent, and individually manageable parts of re-architected, collaborating microservices, with basic capabilities now designated as common shared microservices.

The common services of the DX platform are a foundational part of the unified DX Operational Intelligence (DX OI), which also offers a unified data lake. The DX platform is shared by Broadcom domain monitoring tools and third-party tools, taking advantage of the common services and the data lake for an integrated and unified approach to monitoring.

All services run within a dedicated Kubernetes or OpenShift cluster. Kubernetes and OpenShift are open-source platforms developed by Google and RedHat respectively. These platforms offer efficiency, security, flexibility, agility, scalability, and manageability for orchestrated, collaborating microservices.

Broadcom Solutions and Terms Covered

In this paper, we'll discussing the following solutions and platforms:

- **Application Performance Management** (APM). APM refers to earlier versions of the software, including 10.7 and older—also known as CA APM.
- **DX Application Performance Management** (DX APM). DX APM is the next generation of Broadcom's APM solution.
- **DX Application Performance Management SaaS** (DX APM SaaS). DX APM SaaS is the DX APM solution delivered as a public cloud service.
- **DX Operational Intelligence** (DX OI). DX OI is Broadcom's AIOps platform. This platform consolidates data and events from multiple domain tools, including DX APM, and ingests them into a unified data lake to provide advanced artificial intelligence and machine learning analytics. DX OI on-premises deployments feature these additional components:
  - **DX cluster**. A DX cluster is a customer-provisioned Kubernetes or OpenShift cluster that can be used for the deployment of the DX platform and DX domain monitoring tools. This component is not covered in this white paper—please refer to DX Installer documentation for further information.
  - **DX platform**. This refers to a DX cluster with one or more DX domain tools installed, including the data lake and its common services and features for user management, notification, and integration. This is the foundational platform for DX OI.

Following are a few key characteristics of the new architecture:

- The cloud gateway is now the cluster access service for DX APM agents. The gateway manages agent connections and their Isengard traffic.
- SmartStor has been entirely re-architected, as the network-attached SmartStor (NASS) service based on RocksDB, which is a high-performance, open-source data store developed by Facebook.
- The topology analytics service (TAS) now stores topology data.

- The TraceStore service processes transaction traces, extracting topology data into TAS and forwarding traces to ElasticSearch.
- The ElasticSearch service now stores transaction traces.

These services are highly parallelized, employing streaming, sharding, and rebalancing to provide optimized and load balanced storing and fetching of data for vastly improved efficiency and scalability.

The remaining residual capabilities of the EM that have not yet been re-architected into microservices are present within the DX cluster in these components:

- EM pod collectors provide the residual capabilities as well as transparently delegate processing to cluster services.
- EM pod manager of managers (MoM) of the EM pods.
- For the smallest sizing, a single EM pod runs in standalone mode.

This enables gradual, stable, and transparent transition of residual EM capabilities to re-architected microservices, without breaking any functionality in the process.

This is part of APM's full architectural transition. Over time, the use of residual EM pods will diminish and ultimately the EM pod cluster will vanish. For now, the EM pod cluster needs to be sized as part of the sizing of the DX cluster.

If you need deeper insight into these concepts, please refer to the DX APM Reference Architecture document, which is available through Broadcom's Solution Engineering team.

DX APM Cluster Sizing

With the new architecture, a DX APM cluster's processing and storing capacity is determined by the processing and storing capacity of NASS, TAS, TraceStore, and ElasticSearch services. Capacity is also dictated by the size of the EM pod cluster and the number of EM collector pods, which is controlled by your sizing input.

Rebalancing is a relatively new addition to the microservices' sharding capabilities. This rebalancing eliminates the need to estimate sizing precisely initially. Now, adding NASS and TAS service instances will cause rebalancing—whether to make corrections or accommodate growth—ensuring continuously optimized storing and fetching of data.

Preparation: Sizing

DX APM sizing is based on counts of users, agents, metrics, transaction traces, and vertices. You can use your current cluster's sizing for a full cluster migration or utilize EM and agent size and count metrics for a partial migration.

Within the DX APM cluster, when your tenant is first created, EM pods are either deployed in a standalone fashion or as part of an EM pod cluster. Your sizing determines the number of EM pods deployed. You can have up to 10 EM collector pods per cluster.

The NASS and TAS services have predictable, per-instance capacities for sizing.

Service Instance Capacities

Service	Instance Capacity	Description
NASS	4M metrics	Metric value store
TAS	100K vertices	Topology (vertices, edges) service and store
Cloud Gateway	2K agents	Isengard-aware cloud gateway within DX APM
Cloud Proxy	10K agents	Isengard proxy to DX APM

Multiple Tenants: Capacity and Separation

When using earlier versions of APM, teams could deploy multiple EM clusters, whether for capacity or separation of environments. With DX APM, you use tenants to achieve the same objectives.

Currently, DX APM does not offer any cross-tenant functionality. Given that, in order to use data across multiple tenants, you can use connectors for integration and open APIs for data export. Broadcom's Solution Engineering team can help with these efforts.

Cloud Proxies and Multiple Tenants

Cloud proxies' connections to DX APM are tenant-specific. This means that if you have multiple tenants, you need multiple cloud proxies. In following sections, we provide much more information on cloud proxy options. Specifically, we outline how you can have agents of a cluster redirected to different cloud proxies and can therefore send subsets of data to different tenants, for example for capacity or data alignment.

Best Practice: Multiple Tenants for Environments

It is best practice to have separate tenants for each of your environments, such as development, test, QA, UAT, and production.

Decision: Multiple Tenants for Capacity

If you require DX APM capacity exceeding the capacity of a single tenant, you should carefully consider splitting data across multiple tenants.

The Cloud Gateway and the Cloud Proxy

Key Architectural Changes

When moving to DX APM, your monitoring infrastructure is composed of the following components:

- **Local agents**. Local agents remain deployed as-is, monitoring application servers, infrastructure, and so on.
- **Cloud SaaS**. Your APM EM cluster moves to be part of DX APM. This move can happen at once or be handled partially over time. Optionally, you can retain EM clusters during migration or over time. Connection to DX APM is over the internet for DX APM SaaS, otherwise over your intranet.
- **Optional local cloud proxy**. The optional local cloud proxy enables older agents to connect to DX APM. This gives team ease and flexibility in migration and in fallback. The cloud proxy also enables efficient multiplexing of agent connections through a single DX APM connection.

The diagram below depicts these three parts of the logical DX APM architecture.

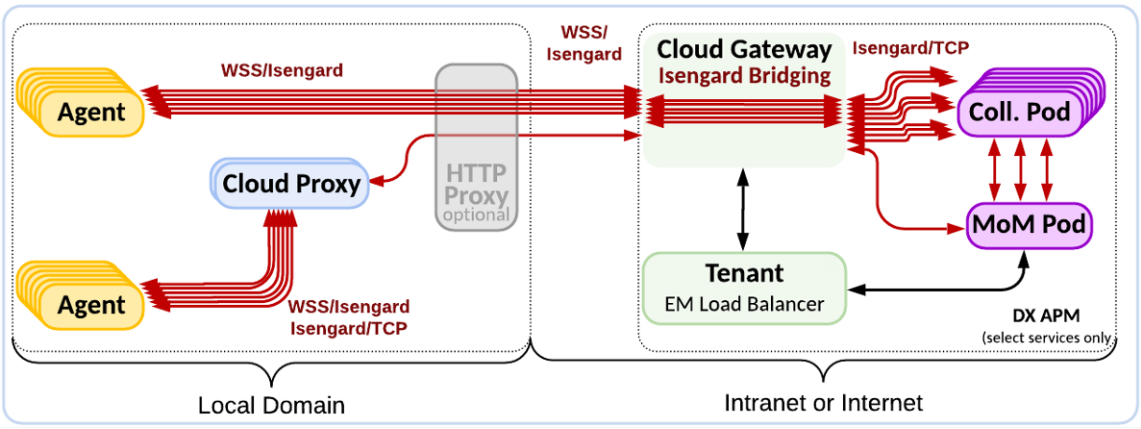

The Cloud Gateway

To access DX APM, agents go through the solution's cloud gateway service. The cloud gateway service only supports WebSocket embedded Isengard. WebSocket is a standardized and efficient Internet protocol for bi-directional communication.

For those agents that only support Isengard over TCP, you can connect to the cloud proxy for proxying to DX APM.

The cloud proxy uses the tenant service and load balancer service to handle authorization, token resolution, connection clamping, and load balancing among EM pods. Therefore, you'll never see reconnection messages in your agent log as that happens transparently within the cluster.

The cloud gateway also extracts and forwards transaction traces to the TraceStore service, and it routes ACC requests to the ACC pod.

There can be multiple cloud gateway instances. This is managed transparently via Kubernetes service routing and handled automatically as per sizing configurations.

The WebSocket Protocol

The WebSocket protocol was adopted because it is an Internet protocol that is standardized, HTTP compliant, and efficient.

The IETF WebSocket Request For Comments (RFC 6433) states that WebSocket "is designed to work over HTTP ports 443 and 80 as well as to support HTTP proxies and intermediaries." This makes WebSocket compatible with HTTP by using an HTTP upgrade header in its handshake to change the protocol from HTTP to WebSocket. This protocol change upgrades a connection that starts as a half-duplex HTTP, into full duplex WebSocket. A full duplex WebSocket connection is essentially a TCP-socket connection that additionally enables the transmission of streams of messages, rather than simply streams of bytes. This includes Isengard messages in particular.

Therefore, WebSocket provides these advantages:

- Allows access through Internet ports 443 for secure connections (and port 80, for insecure connection—not best practice).
- Offers significant stability, performance, and throughput benefits.
- Importantly, WebSocket is designed to be compatible with current HTTP(S) networking security requirements.

Previously, APM did support Isengard tunneling through HTTP(S), but that was abandoned given the above-referenced advantages of WebSocket. With WebSocket and Isengard we can stably connect 10-12,000 agents—through one cloud proxy. By contrast, with HTTP(S) and Isengard, we could barely connect 200 agents. Hence, utilizing WebSocket offers an additional advantage:

- Extremely efficient multiplexing of a large number of agent connections through a single WebSocket connection.

Therefore, by using WebSocket, the required number of connections over the internet to DX APM SaaS may be vastly reduced. Only agents that are APM version 10.7, service pack 2 and above, support WebSocket. However, the cloud proxy can be used to connect agents to DX APM that are version 9.6 or above.

The Cloud Proxy

As its name implies, the cloud proxy acts a proxy or intermediary connection manager between local agents and the cloud gateway. It may optionally be deployed to manage connections to allow these use cases:

- Teams that want improved ease or flexibility in migration and fallback. Further detailed below in In-Place Migration and Selective Migration sections.
- Teams that want to multiplex a large number of agent connections through the single cloud proxy connection to DX APM. This vastly reduces the number of Internet connections used to communicate to DX APM SaaS.
- To enable older agent versions not supporting the WebSocket protocol to connect to DX APM through the cloud proxy:
  - Agent versions APM 10.7 service pack 2 and above support WebSocket/Isengard additionally to Isengard/TCP and may optionally be proxied through the cloud proxy instead of connecting directly to DX APM.
  - Agent versions APM 9.6 thru 10.7 SP2 does not support WebSocket/Isengard and must use Isengard/TCP to connect to the cloud proxy for proxying to DX APM.

The cloud proxy enables ease and flexibility in migration. For some options, fallback is more complex. For all options, fallback will introduce data loss because data stored within DX APM currently cannot be copied to an APM EM cluster for fallback.

A cloud proxy connection to a cloud gateway is tenant specific, so you'll need at least one cloud proxy per tenant.

Multiplexing and High Availability of Cloud Proxies

Customers may ask "If we migrate to DX APM SaaS, will we need to have 10,000 network connections to the cloud, one for each of our agents?"

The answer is: "No, local agents may connect to a handful of cloud proxies, each connecting to DX APM SaaS (possibly via an HTTP proxy of yours) using a single connection."

Multiplexing can be accomplished dynamically or statically in three ways:

- Dynamically, using a single DNS logical host entry for all cloud proxies that resolves to the IP addresses of cloud proxy servers.
- Dynamically, utilizing the virtual IP (VIP) address of a load balancer that is running in front of cloud proxy servers. (In this case, teams would use a single cloud proxy VIP address.)
- Statically, using the MoM's load balancing.xml.

The static option is detailed in the Selective Migration section below.

Best Practice: Dynamic Options for High Availability

The dynamic options are the best practices to accomplish high availability with cloud proxy deployment.

The two dynamic options can take advantage of the cloud proxy's liveness probe (which is found under supportability/health). The liveness probe can be used to exclude non-responsive cloud proxies dynamically. Note that the liveness probe can be used by some DNS implementations to avoid the need for a load balancing device.

The figure below offers a depiction of the two dynamic options.

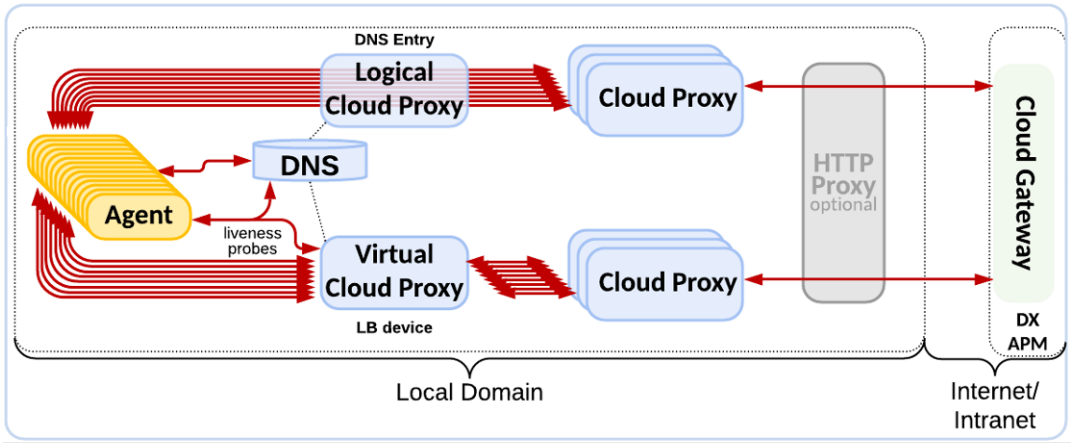

Migration Options and Use Cases

There are three main use cases for the cloud proxy:

- **Option I: In place, at once migration**. This approach entails the immediate migration of an entire cluster. In this scenario, the on-premises EM cluster ceases to exist, and the cloud proxy takes the MoM's place.
- **Option II: Selective migration**. In this option, teams selectively migrate specific agents of a cluster. The on-premises EM cluster continues to operate as needed. Utilizing MoM load balancing, teams can control migration. Teams can also do static multiplexing of agents' DX APM connections via cloud proxies.
- **Option III: Advanced migration**. This approach can be used to address stricter network security requirements. The use of the on-premises EM cluster diminishes over time as cloud proxies are deployed and agents are re-configured. Teams control migration via cloud proxy deployment and agent re-configuration.

Latter sections offer more details on each of these migration approaches.

Option I: In Place At Once Migration

One easy option is to migrate an entire EM cluster in place, at once.

Exchange the MoM for a Global Cloud Proxy

This option exchanges the EM MoM for a cloud proxy, and does so in-place, at once. The EM MoM cluster is stopped, and the cloud proxy starts in the MoM's place.

The following diagram illustrates how a global cloud proxy deployment is established with dual cloud proxies for failover.

Best Practice: Employ ACC

Generally speaking, using ACC, APM Command Center, is a best practice for controlled, tracked, and automated changes to agent configurations. However, for the third migration option above, the use of ACC is specifically recommended. Using ACC will reduce risk in deployment as well as in the case of fallback. Consequently, if you aren't using ACC, deciding on advanced manual migration is the perfect time to carefully consider doing so as multiple changes must be coordinated.

Decision: Migration Approach

Decide on your migration approaches. Mind you, that options may even be combined as well as used selectively per your specific environments and requirements. Following are a couple key questions to consider:

- What are the versions of APM agents that will be migrated?
- Do agent versions and/or migration decisions necessitate deployment of the cloud proxy?

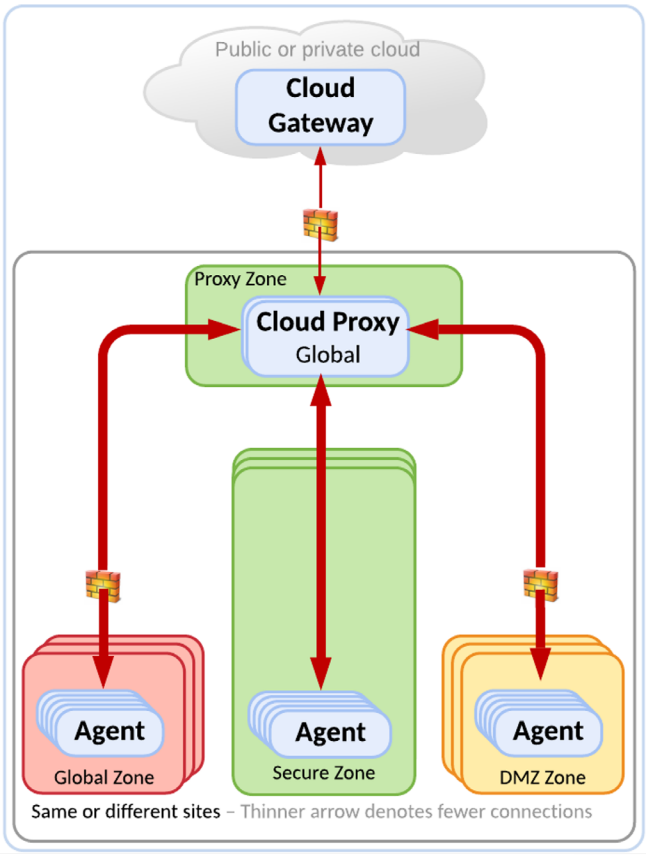

To achieve this, the cloud proxy must be installed on the same server(s) as the MoM. This installation is easy. You can download the cloud proxy from the DX APM tenant download page, which is pre-configured for your tenant connection. Next, you configure your agent channels and eventually add HTTP proxy user/password properties that are identical to the MoM's properties.

Trigger At-Once Migration

To trigger migration, you stop the MoM, start the cloud proxy, and then stop collectors. As the cloud proxy has the same IP address as the MoM and is listening on the same ports, agents will connect to it. As a result, the cloud proxy effectively takes the place of the MoM, without any agent configuration changes.

Fallback

Falling back is as easy as stopping the cloud proxy and restarting the MoM and the collectors—when you have retained the old EM cluster.

Decision and Preparation: Network Matters of WSS, Requirements, and HTTP Proxy

In planning your approach, following are a few networking-related questions and considerations to explore:

- Does your network team have any issues or concerns regarding WebSocket?
- Are there network requirements dictating daisy-chained cloud proxies?
- Do you have HTTP proxies for Internet traffic?
- Cloud proxies must be able to reach DX APM, either directly or indirectly using the cloud proxy.
- An HTTP proxy must allow upgrading to WebSocket in order to pass traffic through.
- HTTP proxy capacity must be adequate for DX APM SaaS internet traffic.

Best Practice: Fallback for In-place, At-once Migration

For eventual fallback teams should retain the old cluster (MoM, collector and database installation(s)) intact until the decision to decommission them is final and you're beyond the point of any need to fall back.

Option II: Selective Migration

Another easy option is to selectively migrate selected agents of a cluster via the MoM over time.

The cloud proxy can't participate in an EM cluster because as a proxy for DX APM it cannot behave as an EM collector. However, you can have an APM MoM "redirect" select agents to connect to the cloud proxy.

Using loadbalancing.xml

Selective migration is done by adding one or more <agent-collector> entries with name="SendToProxy" to your MoM's loadbalacing.xml. These entries match agents by regular expressions of <agent specifier> elements and contain an <include> element for the cloud proxy you want the MoM to redirect matching agents to. The following code sample illustrates how the first SendToProxy entry an agent matches applies to that agent:

```
<agent-collector name="SendToProxy">
    <agent-specifier>.*\|.*\|.*</agent-specifier>
    <agent-specifier>.*\|.*\|.*</agent-specifier>
    <include>
        <collector host="<cloud-proxy-host-1>" port="5001"/>
    </include>
    <include>
        <collector host="<cloud-proxy-host-2>" port="5001"/>
    </include>
</agent-collector>
```

This configures all included cloud proxies as "unmanaged collectors." More on this below.

You may use your own naming for "SendToProxy." For example, you could use "SendToDev" or "SendToUAT" to make the receiving tenant obvious. Alternatively, you can use "CloudProxy1," "CloudProxy2," or "CloudProxy3" to make failover obvious.

Since the release of version 9.1 of APM, agents must connect to a cluster via the MoM to request their allowed collectors' lists. Only the first SendToProxy entry matching an agent will be considered and the MoM will include the cloud proxies of that entry's <include> elements in the returned allowed collectors list. Elements of the list are ordered according to the <include> entries.

Agents will try to connect to cloud proxies in the order of the allowed collectors list until successfully connected. If the agent exhausts the list, a new one will be requested from the MoM.

This causes agents matching an entry to connect to an included cloud proxy, even though the MoM does not control the cloud proxy or include in the cluster. (That's because, as detailed above, cluster EM pods are controlled by the cloud gateway, within the cluster.)

Best Practice: Fallback for Selective Migration

As SendToProxy entries are considered sequentially from the beginning of loadbalancing.xml and only the first matching entry is applied, it is a best practice to add SendToProxy entries at the beginning, so they are considered first. This helps promote clarity and avoid having to change any existing entries. This also makes fall back straightforward; you only need to remove the SendToProxy entries (or make them inactive by treating them as comments).

All agents matching an entry will all connect to the same, first live cloud proxy; they will all receive similar lists to try in identical order.

Should a cloud proxy become non-responsive, its connected agents will simply continue trying its allowed connectors list. Consequently, having multiple <include> elements provide failover.

Selective migration utilizes SendToProxy entries to match agents to be migrated to cloud proxies. This enables teams to add more entries as needed, so they can employ a phased approach over time.

Static Multiplexing with Failover

SendToProxy entries elegantly support another use case: You can do multiplexing of hundreds or thousands of agent connections over a small set of cloud proxies' single connections to DX APM.

The trick is to devise a "SendToProxy" entry that specifies multiplexing. Here's how this works:

- Many agents, hundreds or thousands, may match an entry.
- All agents will connect to the entry's first live cloud proxy.
- A cloud proxy has a single WebSocket/Isengard connection to DX APM.
- This results in the multiplexing of hundreds or thousands of agents over a single WebSocket/Isengard connection.

As mentioned above, this provides failover when there are multiple cloud proxy <include> elements for an entry.

The diagram below illustrates static multiplexing with failover.

Best Practice: Failover for Static Multiplexing

It is best practice to obtain failover by using multiple <include> elements for each SendToProxy entry. For example, if you have three cloud proxies, you use <include> elements in these cyclic permutations for each of the three SendToProxy entries respectively:

- Cloud-Proxy-1, Cloud-Proxy-2, Cloud-Proxy-3,
- Cloud-Proxy-2, Cloud-Proxy-3, Cloud-Proxy-1,
- Cloud-Proxy-3, Cloud-Proxy-1, Cloud-Proxy-2,

Should Cloud-Proxy-1 fail, its agents will connect to Cloud-Proxy-2, and so on. Hence, using cyclic permutations of all cloud proxies adds failover, albeit without load balancing.

There is no load balancing, because failover is static, as configured. The dynamic options provide load balanced failover.

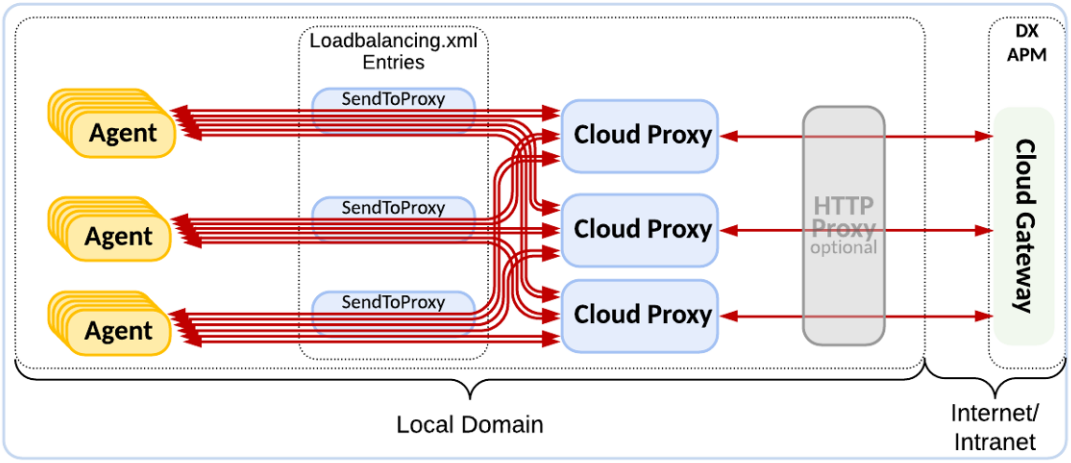

The advantage of this option is the vast reduction in Internet connections obtained, combined with its relative ease of implementation. There's no need for a load-balancing device nor a logical host DNS entry. Effectively, this option handles agents' redirection to a set of cloud proxies solely within your APM configuration.

A Cloud Proxy as an Unmanaged Collector

A cloud proxy can be understood as an "unmanaged collector."

The cloud proxy does have a liveness probe (which appears under supportability/health). However, the MoM does not use this probe to keep track of live cloud proxies. Thus, the allowed collectors list returned to matching agents may include non-responsive cloud proxies. Consequently, the MoM does not hold agents that are to be redirected to a non-responsive cloud proxy in disallowed mode to stall the redirection. For these agents, this has three ramifications:

- The Number Of Disallowed Agents metric does not reflect them. As a result, their existence cannot be alerted on.
- The workstation's status console does not show their non-responsive cloud proxies.
- They will continuously request a new allowed collectors list from the MoM.

Hence, be careful to configure and test SendToProxy entries as misconfiguration and mistakes may indeed cause problems.

For connection issues, always consult agent and cloud proxy logs, rather than the MoM log.

Migrate Gradually Over Time

This scenario allows you to use your existing MoM load balancing to gradually migrate selected agents to DX APM, via the cloud proxy. You only need to change the loadbalancing.xml—requiring minimal effort, with no changes to agents' profiles.

Best Practice: Cloud Proxy Liveness Monitoring

As there aren't any out-of-the-box cloud proxy liveness metrics, it is best practice to use the cloud proxy liveness probe to track these metrics. For example, with the Infrastructure Agent (IA) you can use an EP Agent (EPA) plugin utilizing the liveness probe. (As the successor of the EPA the IA supports EPA-plugins).

Multiple Cloud Proxies

You can even use multiple entries for a number of cloud proxies in order to enable the following efforts:

- Migrating hundreds or even thousands of agents and limiting the number of agent connections to DX APM in favor of fewer cloud proxy connections.
- Migrating agents to multiple tenants—splitting your EM cluster across DX APM tenants. (However, be aware that there are currently no cross-tenant capabilities).
- Migrating agents from multiple clusters into the same DX APM tenant, merging your EM clusters. (However, be aware of any capacity constraints in DX APM.)

Falling back is as easy as removing the SendToProxy entries from loadbalancing.xml.

Changes to loadbalancing.xml do not require a restart as they are immediately applied by the MoM.

Option III: Advanced Migration

Daisy-chaining Cloud Proxies

Cloud proxies can be daisy-chained. Through this approach, one cloud proxy can connect to another cloud proxy, and one cloud proxy can accept connections from multiple cloud proxies. This furnishes teams with indirect options for greater flexibility in enabling agents to reach DX APM through multiple cloud proxies to honor network security requirements. This is detailed further below.

Utilizing Cloud Proxies for Network Segmentation

Deploying multiple cloud proxies allows you to segregate your agents by network zones or segments, without having to change your cross-subnet or zone firewall openings on a per-agent basis. This means you can establish network segregation without having to make a rule change for every new agent; you only need to have one-off rules for cloud proxies. However, agents must be reconfigured to selectively connect to their segment's or zone's cloud proxy.

Per segment/zone cloud proxies connect to global cloud proxies, which are connected to DX APM. This can be achieved either with multiple local cloud proxies or with a single local cloud proxy. The diagram below illustrates these two approaches.

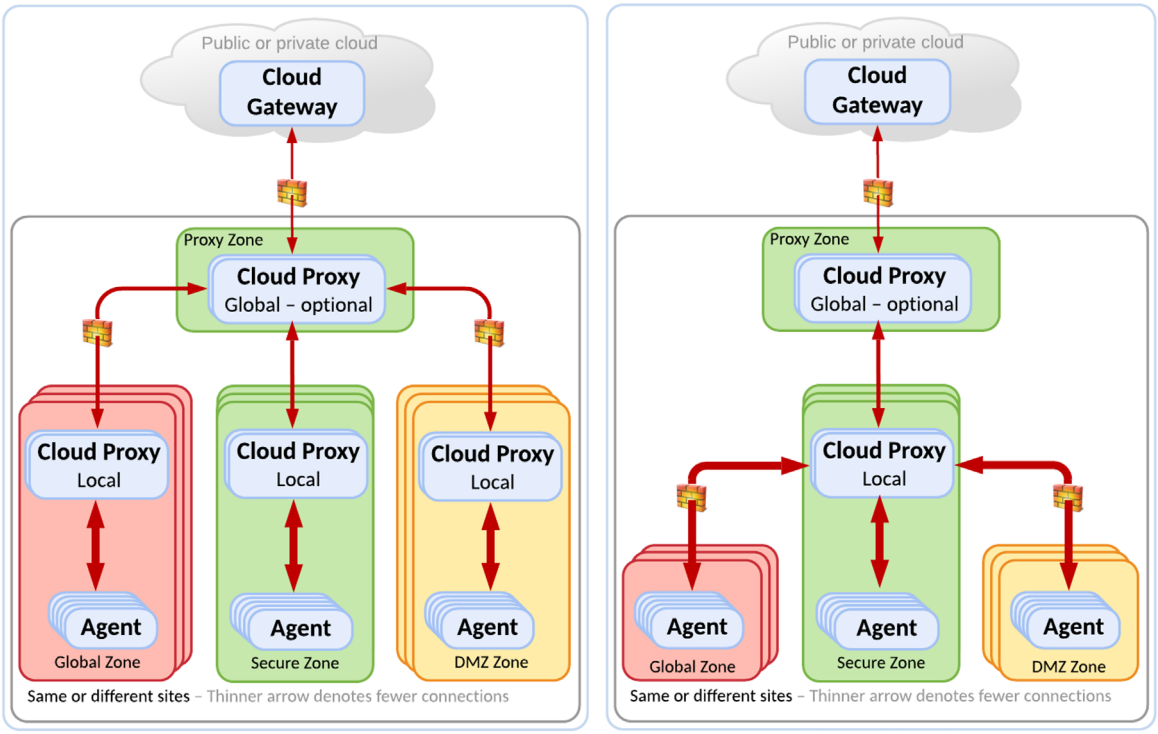

The left option, which has multiple local cloud proxies, could be preferred when one site is involved, that is, the local and global cloud proxies are deployed on the same site and communicate within the site. This enables a more logically separated deployment.

The right option with a single local cloud proxy could be preferred in cases in which multiple sites are involved. For example, if the local and global cloud proxies are deployed on different sites and thus communicate between sites, this can limit inter-site communication as well as number of cloud proxies deployed.

These scenarios reduce complexity and offer a more transparent, efficient, and direct way to address your specific, tightened security requirements within your network topology.

ACC Is Your Friend

This scenario requires changes to agent profiles in order to connect to their segment's cloud proxy. As a result, fallback is more elaborate, thus more prone to error—unless you're using ACC. Thus, only choose this scenario when stricter network security requirements so dictate.

Avoid a Single Point of Failure

High availability configured for your cluster may be carried over for cloud proxies, regardless of your configuration. This requires the following:

- Two agentManager.url properties in agents' profiles for connection to MoMs.
- A DNS logical host MoM entry resolving to the two IP addresses of your MoMs.
- The deployment of two cloud proxies that are reachable at the two IP addresses used for your MoMs.

Fallback

Any data (metrics and transaction traces) moved into DX APM cannot be brought back on to your old APM cluster as no tools exist for that. Thus, if the DX cluster is abandoned, those data are lost.

As your old APM cluster and your new DX APM deployments are completely separate—and retention of the old APM cluster is recommended—fall back is solely about getting agents reconnected to your old APM cluster, away from DX APM or cloud proxies, which is straightforward and fast following migration best practices previously mentioned:

- For in-place, at-one migration fallback entails stopping cloud proxies and starting the MoMs and collectors of the retained old cluster. With no changes, agents will now connect to the old APM cluster.
- For selective migration fallback entails changing back loadbalancing.xml only by deleting or commenting SendToProxy entries added at its top. Within 60 seconds after saving reverse changes agents will connect to your old cluster.
- For the advanced option individual agent profiles must be changed back.

Mind you, that topology will take some time to rebuild but will appear eventually as before, as vertices and edges need to be seen in freshly discovered call paths.

Key Considerations

DX APM Access, Permissions, and SAML

There are two options for managing user access to, and permissions within, DX APM. You can either manually manage users and permissions within the DX APM tenant or use SAML to integrate management of users and permissions, employing your organization's DX Unified Infrastructure Management (UIM) as a SAML identity provider.

Best Practice: SAML

It is best practice to use SAML to always have users and permissions fully aligned with your organization's user management, permissions, and security requirements.

DX APM supports SAML authentication that is initiated by an identity provider. This means that logging on and accessing DX APM starts in your domain, outside DX APM, for forwarding with a SAML payload.

SAML Groups to DX APM Roles

When logging into your domain, users receive their SAML groups, which are assigned by the identity provider. When continuing to DX APM, the SAML groups are mapped to DX APM roles. It is this indirect assignment of DX APM roles from SAML groups that allows a user to log on to DX APM. If a user isn't mapped to any SAML groups, they won't be able to login.

Management Modules

Differences Between CA APM and DX APM

When moving from earlier versions of APM to DX APM, there are differences in Management Module alert notification options.

APM Alerts vs. DX OI Alarms

DX APM is deployed into the DX Platform and therefore shares and utilizes its common features. One such common feature is alerts/alarms. In APM and DX APM, alert is the term used. Within the DX Platform and within DX OI, these are referred to as alarms—for clarity I'll refer to them as DX OI alarms.

This distinction is important to recognize and be clear on. While the terms seem similar, there are real differences. If teams confuse them, it may result in undesired or unexpected behaviors.

The Necessity for Policies

The differences between APM alerts and DX OI alarms stem from the very different heritage and scope of APM and DX OI. Effectively, DX OI has a much broader scope. DX OI is a modern AIOps platform that consolidates data and events from multiple tools and domains and ingests them into its shared data lake. DX OI then offers automated management utilizing that.

For DX OI to support efficient automated management, flexible policies are required. You can't have scalable management without flexible policies. Teams need to have named rules that kick-in based on general logical conditions associated with processed data. Thus, DX OI has added policies for alarms for manageability. On the other hand, APM with its much narrower scope, can have locally determined, simple conditions for real-time processing.

There are two best practices use cases that help teams bridge alerts and alarms meaningfully and usefully. One use case is for DX APM direct alerts, and one is for utilizing DX OI policy alarms.

Also, we have created the EasyMigrator tool for automated migration of Management Modules, domains, alert email actions, and more. Assessing the feasibility of these tools should be your first item on your migration agenda. (See the EasyMigrator Tool section below.)

Decision: SAML

Use SAML for user and permission management, and, if that is not an option, establish processes for manual user and permission alignment.

Preparation: SAML

It is recommended that you involve your SAML peer to configure URLs and map attributes and groups between SAML and DX APM. This collaboration is often needed for ensuring a smooth migration. The tricky parts of configuration and integration typically happen on the SAML identity provider side and require adequate permissions.

Preparation: Assessing EasyMigrator

Please assess carefully if EasyMigrator is feasible for automating your migration of Management Modules.

If not, please consider reaching out to Broadcom's Solution Engineering as we'd like to fully understand your challenges and possibly extend EasyMigrator for you.

Notifications

For defining new alerts and notifications, you need to know how DX APM notifications are supported within the DX platform.

Channels

One DX Platform concept you need to understand for both use cases is channels. A channel is simply a named mechanism for sending a notification from a sender, in this case DX APM, to some recipient, for example an email server—in which case we refer to it as an email channel. Other channels might send notifications to Slack, a service management solution, or a webhook endpoint.

Message Template

Another important concept to understand is the message template. As the name suggests, it is a template for the contents and layout of the notification email sent.

A message template can be shared and assigned to one or more email channels. Hence, you can start with a few and expand as you go—having your naming standard in place.

APM Direct Alerts

If your alerts have the same action for caution and danger, you can use APM direct alerts. To use this, you simply create the channels the notification uses to reach the recipients.

Do not attach any policies to channels because if you do, the policies will kick-in and process APM alerts as DX platform alarms, in addition to APM's alert processing, which will most likely cause duplicate notifications.

Decision: Necessity for DX OI Policy Alarms

If your alerts have different actions for caution and danger, you must use DX OI policy alarms. If not, use DX APM direct alerting.

Preparation: Integration Parameters

If you plan to use other channels than email alone, you must gather integration parameters.

Best Practices: Reusable Channels and Alerts for Manageability and Consistency

For best practices there are two things that should be considered beforehand:

- **Needed channels**: For optimal reuse of channels for alerts, you should determine the set of your different recipients of alerts. Include team, project, owner, line of business, and so on in your considerations to reach a set of recipients in which each recipient has the same management or lifecycle scope. And be mindful that for email channels, the email addresses are part of the channel's definition because you probably want to limit updating channels simply to add or remove those. Establish a naming standard covering this and then define channels. This should give you the least number of consistently named channels that are most easy to recognize and then use those intentionally when attaching channels to policies.
- **Alert naming**: For optimal use of channel names in policy conditions, it is best practice to adopt a naming standard for channel names. You should use the same considerations as for channel names above. Specifically, this will make it easy and transparent to limit policies to apply for intended teams, projects, lines of business, and so on, using the contains operator in conditions. (There are plans to propagate APM attributes to alerts so they will be usable for policy conditions, which will offer a better option than only relying on naming standards for policy conditions.)

Then, use the defined channels for your alerts when defining them using the very same features available in APM 10.7. As alerts defined using channels bypass DX platform alarm processing, nothing further is required, which is why I refer to this as APM direct alerting.

A channel can be used for as many alerts as you see fit, which should be utilized to obtain a best practice limited set of channels.

DX Platform Policy Alarms

If your alerts have different actions for caution and danger, you must use DX platform policy alarms. This is because, for simplicity, in DX APM an alert can only have one channel used for both caution and danger.

To utilize DX platform alarms, you need to know that the DX platform receives and consolidates all alerts from domain tools, including APM. This means that all APM alerts are always continuously passed to DX platform alarm processing.

Every alarm is processed with respect to every policy that has a channel attached. A policy is a set of logical conditions evaluated on alarm attributes. Whenever an alarm meets a policy's conditions, it is sent through the channels that are attached to the policy.

Hence, to use different channels for caution and danger, you must define two policies, one that evaluates to true for caution alarms and another one that evaluates to true for danger alarms. That is done using the alarm conditions Severity=Major and Severity=Critical, respectively. (Note, APM caution severity equals DX platform major severity).

When using DX platform policies for alarms, do NOT specify any channel for the alerts in DX APM. Because, as noted, that will additionally cause DX APM direct alerting through that channel, likely duplicating your notification.

Best Practice: Reusable Policies for Manageability and Consistency

It is best practice to adopt a standard for policy names to best guarantee keeping separate things separate, that is, only attaching intended channels to policies. You use the same considerations as for channel names above. EasyMigrator uses this option—even when caution and danger alerts have identical email actions—as is generally the case.

Custom Dashboards

We are in the process of developing a tool that will offer semi-automated migration of custom APM dashboards to DX dashboards.

However, options in the two are quite different. As a result, doing a full migration and keeping all features is unfortunately not possible, neither manually nor automatically. The automated dashboard migration is intended to provide you with the best possible starting point for creating DX dashboards based on the data content of your APM custom dashboards. Broadcom's Solution Engineering team would be available to initiate a trial for migrating your custom dashboards.

EasyMigrator Tool in the EasySeries

Over the years many customers have developed numerous Management Modules, email alerts, domains, and more, that need to be carried forward when migrating to DX APM.

Additionally, teams have to map each and every domain to a universe, associate Management Modules with a universe and create experience views and email channels. For some customers, there could be hundreds of these elements, which could take a very long time to migrate manually.

With the EasyMigrator tool you can do almost all these tasks in minutes and only worry about customization and one-off configuration items. Moreover, with EasyMigrator's simple configuration, you can map multiple APM clusters to your DX APM tenants and run it once for all migrations. Value from your DX APM migration will be obtained almost immediately.

Please, consider and investigate thoroughly if this option is viable for you. If not, you are urged to reach out to Broadcom's Solution Engineering team to discuss your challenges. We may be able to expand or adapt EasyMigrator to address your specific requirements. Please know that we are eager to support automated migration using EasyMigrator.

The Gist of EasyMigrator

The gist of EasyMigrator is to migrate:

- Management Modules, including metric groups and alerts.
- Domains as domain attributes—as per best practices.
- Domains into universes with access, experience views and service definitions.
- Email actions into email notifications creating necessary channels, message templates, and policies.
- Not included:
  - Data is not migrated.
  - Custom dashboards are not migrated. Please see the Custom Dashboards section for more information.

What EasyMigrator Does for You

For you to hopefully appreciate the value of EasyMigrator and what exactly it does for you, here's a detailed look at the two phases the tool executes:

1. Migrates Management Modules from CA APM to DX APM as-is, then analyzes and migrates email actions.

Differences:

- Actions do not exist for DX APM alerts.
- Channels, message templates, and policies are new in DX APM.

Consequential limitations:

- Only email actions can and will be migrated.
- Other actions require a local processing environment and cannot be migrated.

Migration:

- Metric groups and alerts are created
- Since universes are based on attributes rules. Agents that do not possess any attribute may need to be included manually.
- Alerts having email actions channels, message template, and policies will be created: Two policies, for caution and danger respectively. A common message template, as per the Management Module, that includes:

  - Header: APM [Severity] Alert for [Management Module] – [Alert Status].
  - Body: Template layout containing alarm attributes.
  - Isolation View Link, which is dynamically generated by EasyMigrator. This link provides navigation into a Team Center-focused isolation view with identified anomalies or problems, enabling direct root cause analysis. (A currently private link endpoint is used.)
  - Alarm URL: Using ${alarmURL} substitution variable.
- One email channel per email action associated with the two policies and the Management Module message template. Email recipients are migrated.

2. Migrates domains. In essence, domains are migrated into universes:

Domains of domains.xml file are traversed and universes reflecting them created:

- Creates domain name attribute rules by using the domain agent expression.
- Creates an experience view and a DX OI service using the domain attribute.
- Associate the domain's Management Modules with the universe.

We strongly encourage that teams test EasyMigrator within a non-production implementation first.

Preparation: EasyMigrator Connectivity

For EasyMigrator, you must collect the agent connection URL and access token of your DX APM tenant and CA APM cluster configuration.

Customer Example

Broadcom's Solution Engineering team recently worked with a customer in France. Their APM team wanted to prepare for a proof of concept (POC) for DX APM.

Initially, the customer's team decided to deploy DX APM in an on-premises Kubernetes cluster. One reason is that they wanted to ensure they adhered to the company's policies pertaining to the General Data Protection Regulation (GDPR). However, we managed to persuade the customer to do a POC using the DX APM SaaS solution—so they could avoid having to do on-premises deployment of the DX cluster.

The team chose deployment option I. This option exchanges the EM MoM for a cloud proxy—and does so in-place, at once. When the cluster is stopped, the cloud proxy is started in the MoM's place. For this customer, we had 1,000 agents connected in about one hour.

To achieve this, the cloud proxy was installed prior to our POC visit by the customer on the same server as the MoM. This installation is easy. You download the cloud proxy from the DX APM tenant download page, pre-configured for your DX APM tenant connection. Next, you configure your agent channels and eventually add HTTP proxy user/password properties, which are identical to the MoM's properties.

The customer's team was very pleased with the ease of migration from CA APM to DX APM—in this case the public cloud SaaS option. As a result, they've since decided to go into production with DX APM SaaS. Further, they've been able to make this move, while ensuring compliance with GDPR and other relevant policies.

Conclusion

For teams who are planning or considering a move to DX APM, there are a range of options and factors to consider. Our hope is that this white paper will provide useful insights in making this move most efficiently, and ultimately realizing maximum value from the next-generation DX APM solution.

About the Author

Henrik has worked with Broadcom and CA Technologies for the last 23 years, and for the past 15 years has been working on APM solutions. His experience spans DB2, data modeling, portals, security, and APM. Henrik has worked within presales, customer lifecycle solutions, SWAT, and solution engineering teams. An expert in helping customer teams master APM value, Henrik operates with a laser-like focus on increasing customer knowledge and satisfaction. Henrik holds a master's degree in Computer Science from the University of Copenhagen and has 44 years of experience in the IT industry.

About Broadcom

Broadcom Inc. (NASDAQ: AVGO) is a global technology leader that designs, develops and supplies a broad range of semiconductor and infrastructure software solutions. Broadcom's category-leading product portfolio serves critical markets including data center, networking, enterprise software, broadband, wireless, storage and industrial. Our solutions include data center networking and storage, enterprise and mainframe software focused on automation, monitoring and security, smartphone components, telecoms and factory automation. For more information, go to www.broadcom.com.

1221

BROADCOM SOFTWARE