BAYSHORE NETWORKS
INDUSTRIAL AND IT NETWORK SECURITY

INFO@BAYSHORENETWORKS.COM
WWW.BAYSHORENETWORKS.COM

A NEW PARADIGM FOR NETWORK SEGMENTATION

@BAYSHORENET

Introduction

It seems like we've been hearing about the virtues of network segmentation for a very very long time. Every month, however, we find ourselves at customer sites where their implementation of the technique is no better than 'partial'. It's not a question of failing to understand the goal, or not appreciating the benefits; instead, it's a problem of bridging the gap between where they are today, and what's required to implement a thorough and sustainable segmented network in the future. It is, as the saying goes, a bridge too far.

The reasons are well understood, and often directly anticipated by other white papers on the subject. Most solutions take a whole-network approach to planning a solution. Our friends at GE Digital published a document back in 2015 talking about the importance of virtual segmentation strategies and zones . It's a solid paper, with great visuals, and it highlights a few important points.

- *Industrial assets are often in fixed locations and cannot be moved*
- *True physical network segmentation is therefore even more complex once you think about connecting all those physically distant and separate endpoints*
- *Traditional enterprise network techniques, such as VLANs and routing rules, are difficult to manage with the limited staffing of typical ICS networks*

The solution – virtual segments and zones – has been adopted by a number of different network and security vendors. The idea is simple: define assets into logical groups, regardless of their real locations, and apply traffic-permission and content-enforcement rules per group. If you want to define a list of devices and ensure that they will only communicate via EtherNet/IP, you can do that so long as all the control points respond to the grouping policies.

Such approaches, however, do still depend on a whole-network scope of visibility. Somewhere there must be a device that physically inspects all the traffic and makes the determination of access and execution. Dealing with exceptions and corner cases is therefore dependent on managing that central control, potentially upending your hard work if you get it wrong and don't notice.

At Bayshore we've decided on a simpler approach. Why force the customer to segment the whole network at once? What if, despite accepting that their existing production/ICS/critical infrastructure networks are essentially flat and unsegmented, they could carve out individual segments at will and then apply very rigorous controls against activity which somehow touched them?

This white paper will explore that question. We'll evaluate the comparison of whole-network virtual segmentation versus per-asset microsegmentation, and offer some data points on relative cost, relative strength of security controls, and ease of implementation.

Whole-Network Segmentation Approaches

Large and sophisticated IT departments often run comprehensive segmentation strategies on corporate or enterprise networks. They govern what types of traffic can move within and between zones, and are typically organized by business function. One segment might host application servers, another might have back-end databases, a third might be user desktops within a particular facility. The options are endless. Most of the time, some amount of interaction between segments is necessary, and there is always some room for debate about how valuable those delineations are if intruders or malware can find a way to propagate between them. Think about the number of ransomware stories that begin with someone with a standard personal workstation who visits the wrong web page (via permitted internet access) or downloads the wrong email attachment. Damage is rarely limited to just the adjacent personal workstations.

Nevertheless, for our discussion of production, ICS, critical infrastructure networks, we're going to accept that the majority of asset owners have neither the human capital nor the budgets to build and maintain such a system. As a result their options skew more towards "soft" segmentation strategies. We see these defined towards a layer 3 approach, using subnets, more often than we see true layer 2 VLANs, because subnets are easier to manage. They're not as effective as VLANs at constraining traffic, however, primarily because VLANs create separate broadcast domains at the data link layer. This means a device on one VLAN cannot find a device on another VLAN unless there is a layer-3 router in place as well as the switch administering the VLANs themselves.

A subnet, in contrast, is less about security separation than it is about planning for network growth. By creating multiple Host IDs within your primary Network ID, you have more room to expand as you add devices to the network. This is good in rapidly-evolving environments, but many production/ICS/critical infrastructure tend not to grow that quickly. Most customers end up using subnets for easy visual identification of functions within their plant: one Host ID per building, or one per physical site. The small team that runs the network learns very quickly that anything on the 10.x network is building 1, 20.x is building 2, and so on. It's a work-saving shortcut rather than a meaningful networking advantage.

In short, VLANs can work well if the customer has the necessary switching and routing equipment, and the discipline to study all their assets and network traffic to build meaningful segmentation. Many smaller asset owners may not find the benefits to be worth the effort, and as a result tend to default to subnets for simplified workflow while avoiding significant security benefits.

Targeted Segmentation Approaches

If we return to the question posed earlier, what does it mean to carve out an individual segment? For our purposes we're assuming the segment definition itself is meaningful, i.e. that it enables some kind of security advantage. Our goal is to give the asset owner additional control over traffic within that one segment, without having to do a whole assessment and network redesign and accept the ongoing maintenance burden that comes with it.

For our purposes we're going to explore two types of targeted segmentation, which we'll refer to as

- *Endpoint Segments*
- *Trusted Domain Segments*

What is an ‘Endpoint Segment’?

An Endpoint Segment is a small number of assets (typically 10 or fewer) of devices, likely all within a single cabinet or at least a single building, which collectively perform one basic function within (drop the) production/ICS/critical infrastructure networks. It might be a VFD, PLC, and HMI that collectively run a lift pump at a water plant, for example, or a skid in a manufacturing environment. In general most asset owners will agree that the scope of work that Endpoint Segment is expected to perform is both very well-understood and changes infrequently if at all.

It is therefore possible to define enforcement parameters around network traffic which enters or exits that Endpoint Segment. Most, or all, of the following statements will generally be true

- *The assets outside the Endpoint Segment which need to interact with assets within it are all known*
- *The traffic patterns between those assets and the Endpoint Segment assets are easily observed and don’t fluctuate*
- *Exceptions are rare and usually can be anticipated in advance, such that scheduled changes may be set up in advance*

We want to define enforcement around these parameters because it allows us to extend security protection down to the assets in that Endpoint Segment. With the right tools, we can build the rules necessary automatically, and there is a low practical risk of false positives because of the relatively static nature of assets and network communications within an ICS environment.

BEFORE

Typical Flat Network Architecture

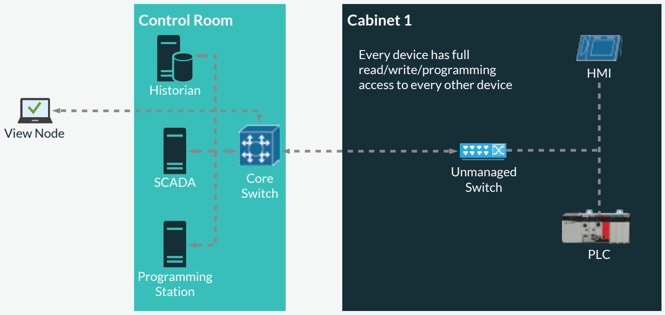

Every device has full read/write/programming access to every other device

AFTER

OTfuse Use Case - Endpoint Segment

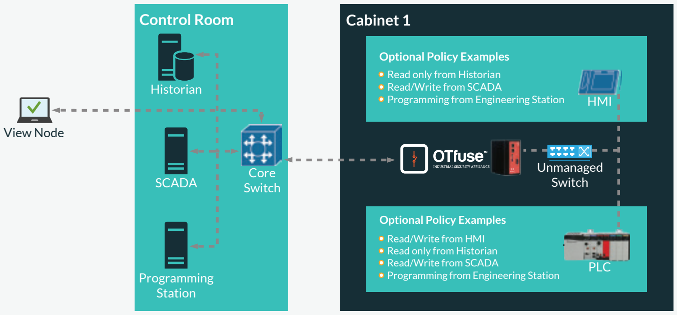

Optional Policy Examples

- Read only from Historian
- Read/Write from SCADA
- Programming from Engineering Station

Optional Policy Examples

- Read/Write from HMI
- Read only from Historian
- Read/Write from SCADA
- Programming from Engineering Station

Bayshore offers a family of automated security appliances –OTfuse™ – which do exactly that. They work as a layer 2 switch (or, more precisely, a transparent bridge) which sits in front of the cabinet, taking network traffic that is going to and from the assets inside, directly off a network switch. It learns all the source and destination IP addresses for all network flows, in either direction, and classifies the traffic according to the port and industrial protocol used. It further analyzes all the protocol activity using a Deep Content Inspection approach, to build rules which are organized into three classes of activity types:

- *Read only*
- *Read & write*
- *Programming update / full access*

OTfuse Customer Deployment

The customer needs only activate the OTfuse in learning mode, let it build its rules, and then review the suggested policies. Any required changes can be made manually, and tested in production by switching the device into 'Monitoring' mode. In this configuration, violations of policy will be recorded, but no action is taken to constrain or filter out traffic. Customers can add, with one or two clicks, the anomalies which they actually want reflected in policy going forward, and the others will remain excluded.

Once switched into 'Protection' mode, the OTfuse will apply its customized policy to the Endpoint Segment assets directly and in real time, while alerting the customer via their existing SCADA system of any blocked events. Customers can interact with a single OTfuse directly from their control room, or make use of a centralized web-based management console for batch review of policies and events across a large number of individual appliances.

Bayshore also offers variants of OTfuse which are designed to create an Endpoint Segment around critical SCADA application servers themselves. The first of these is built for GE Digital's iFIX 6.x family of SCADA solutions, and it works natively with the proprietary protocols iFIX uses to communicate between the SCADA servers, the various types of view nodes, and the adapters, collectors, and drivers required to interact with automation devices elsewhere in the plant.

Ultimately the advantages of using OTfuse to define and enforce Endpoint Segments boil down to faster implementation, lower cost, and more flexible control, all compared to the legacy approaches using firewalls, switches, and routers. Even if the control network is left entirely unchanged except for one Endpoint Segment defined by one OTfuse for the most critical asset, the asset owner will still realize a significant improvement in protection around that device. It can be implemented in an afternoon and the costing is a small fraction of any other kind of partial- or whole-network redesign effort.

What is a 'Trusted Domain Segment'?

The crucial distinction between a Trusted Domain Segment and an Endpoint Segment concerns the prevalence of traffic activity. In an Endpoint Segment, traffic on multiple ports and protocols is expected to show both ingress and egress patterns, where connections are initiated either within the Endpoint Segment, or outside of it.

In a Trusted Domain Segment, we generally only want traffic to cross the boundary point when it is explicitly initiated by devices within the Trusted Domain. In other words, we expect a preponderance of unidirectional activity from the trusted domain to an external domain. Under some circumstances, we may also allow responses, but only if they are statefully related to a request triggered by the trusted domain in the first place and originate from pre-defined sources and content types.

In general, Trusted Domains report data about their activities out to management systems or other stakeholders, but have no regular need to receive instructions or updated programming. These systems are the most stable, the most critical if attacked or subject to unauthorized use, and ultimately, the systems on which the asset owner places the highest commercial and operational importance.

A Trusted Domain Segment should therefore be protected by the most rigorous forms of security, regardless of standards or technologies in use elsewhere in the production/ICS/critical infrastructure networks. The most effective method to protect such a Segment from network attacks is, logically enough, to prevent those attacks from being able to connect into the Trusted Domain in the first place.

In the past, many Industrial networks were "air gapped" from the wider world. An air gap is just that – a physical break in space without any kind of network connection – but over time most whole-network airgaps were eroded or compromised by demands for more flexible usage. The interoperability demands between ICS networks and upstream stakeholders, coupled with the need, even if occasional, for internet connectivity, makes it very difficult to insulate an entire large network via an air gap. Most assessments of such networks find some kind of back channel or workaround, which renders the attempt largely meaningless.

It is possible, however, to use similar isolation techniques on a much more targeted scale. If you stand 20 yards away and look at a tree full of leaves, it's not easy to trace exactly which branch any given leaf or twig is attached to, nor to trace the full and accurate path to the main trunk. But if you go up to the tree and hold a single branch in your hand, it's much easier to have absolute certainty on those connection paths.

The same is true with Trusted Domain Segments. We focus only on very small asset counts within such a grouping, and we enumerate every possible connectivity requirement, with no gray areas or ambiguities. As a policy, we accept that there can be no slippage in these definitions.

And we can protect such a group via network isolation techniques. Our goal is to ensure that, when needed, the assets in the Trusted Domain Segment can communicate upstream with the full benefit of speed, low latency, and functionality, as if the network were not isolated at all. But under testing conditions, we can demonstrate with certainty that no return path networking is possible.

As with Endpoint Segments, Bayshore offers a solution to create and protect a Trusted Domain Segment. We call this solution NetWall™ and it is intended to permit unidirectional flows from a trusted domain to an external domain, under only explicitly-defined conditions.

The way it works is by using two separate servers. Each one has a standard ethernet interface. One is connected exclusively to the trusted domain, the other is connected exclusively to the external domain.

Between the two servers is a PCI Express interface cable. The trusted side is able to deliver data to the untrusted side in accordance with rules set up by the administrator. These rules all include at least the following.

- ***Permitted source IP address & port on trusted domain***
- ***Intended destination IP address & port on untrusted domain***
- ***A rule type, which can be TCP, UDP, file transfer, OPC, or Modbus/TCP***
  - *For File Transfer, OPC, and Modbus/TCP, various detailed parameters are required to configure the full connection*

Simplified Use Cases for NetWall

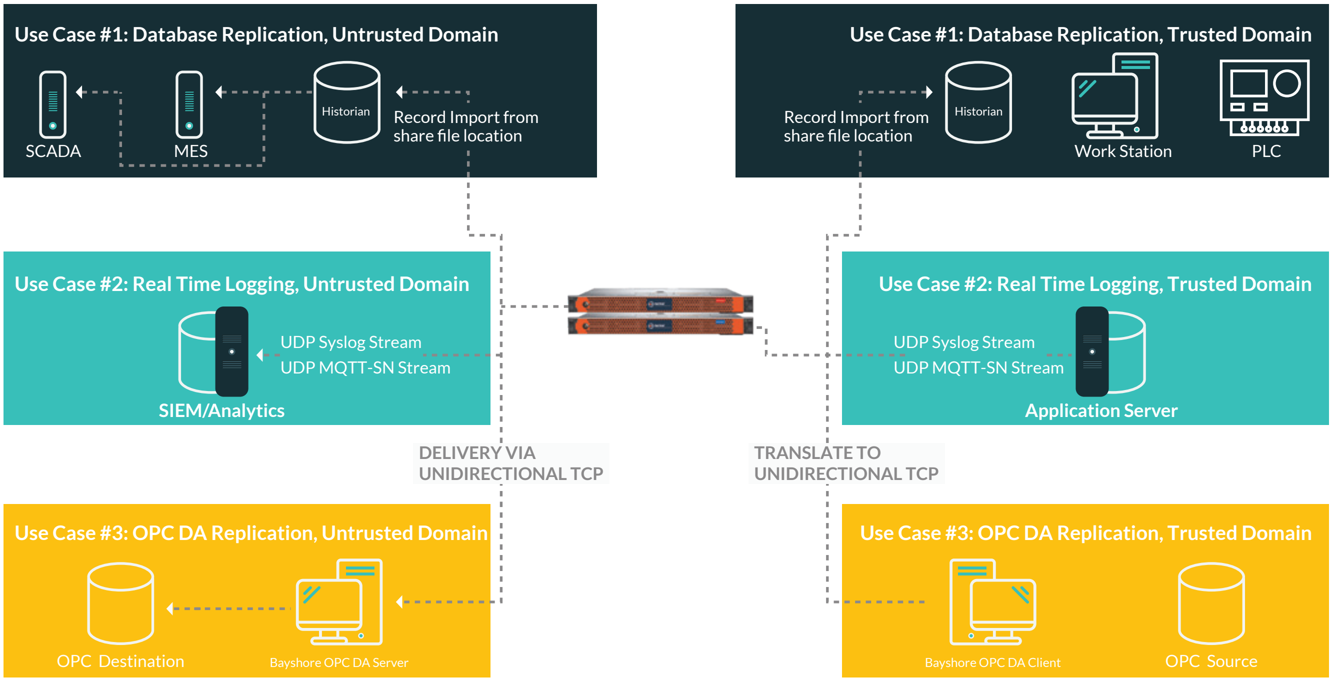
Use Case #1: Database Replication, Untrusted Domain
SCADA
MES
Historian
Record Import from share file location
Use Case #1: Database Replication, Trusted Domain
Record Import from share file location
Historian
Work Station
PLC
Use Case #2: Real Time Logging, Untrusted Domain
UDP Syslog Stream
UDP MQTT-SN Stream
SIEM/Analytics
Use Case #2: Real Time Logging, Trusted Domain
UDP Syslog Stream
UDP MQTT-SN Stream
Application Server
DELIVERY VIA UNIDIRECTIONAL TCP
TRANSLATE TO UNIDIRECTIONAL TCP
Use Case #3: OPC DA Replication, Untrusted Domain
OPC Destination
Bayshore OPC DA Server
Use Case #3: OPC DA Replication, Trusted Domain
Bayshore OPC DA Client
OPC Source

Once configured for functionality other than file transfers, the trusted domain server will listen via filtered ports on its local network interface for incoming connections from the permitted source(s). As activity is received, it is disassembled, the payload and connection metadata are delivered to the second server on the untrusted side, and it in turn reassembles the data into a new network connection delivered only to the permitted destination.

The trusted domain server can guarantee that messages have been successfully delivered to the external domain because, before transmission, it computes a verifiable checksum for the message it is about to transmit.

The internal bandwidth between these servers is very high – up to 40 gigabits per second – so there is ample headroom for multiple concurrent delivery sessions. NetWall is available in various performance tiers up to and including gigabit speed, and we have the option of easily building multi-gigabit versions in the future.

SUMMARY

The benefits from deploying NetWall for your network segmentation needs for Endpoint Segments and Trusted Domain Segments are very straightforward:

- *Isolate the Trusted Domain Segment with a verifiable electronic security perimeter*
- *Ensure that malware or unauthorized probing and reconnaissance efforts have no chance of entering the Segment*
- *Provide full-bandwidth capabilities to replicate data into the external domain*
- *Guaranteed delivery of your data*
- *One platform grows with your needs from 50 Megabit/sec to 10 Gigabit/sec through field upgrade software license keys*

BAYSHORE NETWORKS
INDUSTRIAL AND IT NETWORK SECURITY

CORPORATE HEADQUARTERS

4625 CREEKSTONE DR, #100
DURHAM, NC 27703

SALES AND PARTNERSHIPS

TOLL FREE: (844) 200-7181
EMAIL: SALES@BAYSHORENETWORKS.COM

TECHNICAL SUPPORT

TOLL FREE: (844) 200-6546
EMAIL: SUPPORT@BAYSHORENETWORKS.COM